

Surface-Electrode Ion Traps for Scalable Quantum Computing

David T. C. Allcock

A thesis submitted for the degree of Doctor of Philosophy

Hertford College
Oxford

Michaelmas Term, 2011

Abstract

Surface-Electrode Ion Traps for Scalable Quantum Computing

A thesis submitted for the degree of Doctor of Philosophy

Michaelmas Term 2011

David T. C. Allcock

Hertford College, Oxford

The major challenges in trapped-ion quantum computation are to scale up few-ion experiments to many qubits and to improve control techniques so that quantum logic gates can be carried out with higher fidelities. This thesis reports experimental progress in both of these areas. In the early part of the thesis we describe the fabrication of a surface-electrode ion trap, the development of the apparatus and techniques required to operate it and the successful trapping of $^{40}\text{Ca}^+$ ions. Notably we developed methods to control the orientation of the principal axes and to minimise ion micromotion. We propose a repumping scheme that simplifies heating rate measurements for ions with low-lying D levels, and use it to characterise the electric field noise in the trap.

Surface-electrode traps are important because they offer a route to dense integration of electronic and optical control elements using existing microfabrication technology. We explore this scaling route by testing a series of three traps that were microfabricated at Sandia National Laboratories. Investigations of micromotion and charging of the surface by laser beams were carried out and improvements to future traps are suggested. Using one of these traps we also investigated anomalous electrical noise from the electrode surfaces and discovered that it can be reduced by cleaning with a pulsed laser. A factor of two decrease was observed; this represents the first *in situ* removal of this noise source, an important step towards higher gate fidelities.

In the second half of the thesis we describe the design and construction of an experiment for the purpose of replacing laser-driven multi-qubit quantum logic gates with microwave-driven ones. We investigate magnetic-field-independent hyperfine qubits in $^{43}\text{Ca}^+$ as suitable qubits for this scheme. We make a design study of how best to integrate an ion trap with the microwave conductors required to implement the gate and propose a novel integrated resonant structure. The trap was fabricated and ions were successfully loaded. Single-qubit experiments show that the microwave fields above the trap are in excellent agreement with software simulations. There are good prospects for demonstrating a multi-qubit gate in the near future. We conclude by discussing the possibilities for larger-scale quantum computation by combining microfabricated traps and microwave control.

Dedicated to the memory of Tom Goodwin.
This is all your fault you know.

Acknowledgements

I have thoroughly enjoyed working with my supervisor David Lucas on this thesis. The variety of experiments presented here is testament to the encouragement and support he has given me in pursuing new avenues of research. He provided the attention to detail that ensured projects were feasible before being embarked upon and scientific corners weren't being cut. I'm also very grateful for the patience he's shown in making good some of the rather terrible first drafts I've handed him over the years.

Jeff Sherman was my post-doc during the early part of this project and gave me the all-important 'lab education' that got me started in the ion trapping game. Tom Harty built up most of the microwave experiment and without his input the designs would still be firmly on the drawing board. Luca Guidoni took time off from his own group at Paris Diderot to work on the laser cleaning project, to which he added much expertise. Derek Stacey guided me through much of the atomic physics in this thesis and his mastery of the Bloch equations was a constant asset to these experiments. Andrew Steane laid the foundations of ion trapping at Oxford and offered much useful advice. Paul Ewart and Ben Williams kindly lent us one of their group's pulsed lasers for the laser cleaning chapter. Simon Webster was the bestest office mate ever.

This thesis wouldn't have been possible without the work of all the other 'new lab' students and post-docs over the last six years- Michael Curtis, Gergely Imreh, Hugo Janacek, Ben Keitch, Diana Prado Lopes Aude, Eoin Phillips and Nick Thomas. The advice and company of the 'old lab' crew, Chris Ballance, Alice Burrell, Jonathan Home, Norbert Linke, David Szwerc and Simon Webster has also been invaluable. Thanks also go to the technical and administrative staff in the department, particularly Keith Belcher, Pete Clack, Sue Gardner, Rob Harris, Rick Makin, Simon Moulder, Paul Pattinson and Graham Quelch.

Matt Blain, Mike Descour, Dave Moehring, Dan Stick and the rest of the Sandia team have been invaluable collaborators during this thesis, both in terms of their scientific input and their awe-inspiring trap fabrication efforts. Thanks go to Prof. Rainer Blatt, Prof. Piet Schmidt and the Innsbruck group for their hospitality during the summer I spent there at the start of my DPhil. Members of Dr. Dave Wineland's group at NIST have been generous with their encouragement and advice on the microwave project.

I gratefully acknowledge the financial support of EPSRC and Hertford College's Carreras senior scholarship.

Thanks to all my friends I've made in Oxford for all the funtimes, especially the IMSoc kids, the Yacht Club, my flatmate Alex and the various Linacreites.

Vanessa and my family have been great at putting up with me and supporting me over the years. Thanks especially to my Mum for always being ready to help when I needed to move house or I ran out of 'beer tokens'.

Contents

Contents	vii
1 Introduction	1
1.1 Ion trap quantum information processing	2
1.2 Focus of the thesis	6
1.2.1 Improving fidelity	7
1.2.2 Improving scalability	10
1.3 Structure of the thesis	11
2 Experimental apparatus	13
2.1 Laser systems	13
2.1.1 Photoionization lasers	14
2.1.2 Ca ⁺ lasers	14
2.1.3 Laser diagnostic system	18
2.1.4 Trap optics	20
2.2 Experimental control	24
2.2.1 Experimental control computer	24
2.2.2 Auxiliary computer	25
2.2.3 Fibre noise-eating FPGA system	27
2.3 Imaging system	28
2.3.1 Filters	29
2.3.2 Imaging computer	30
2.4 Trap fabrication	30
2.5 Vacuum system	32
2.5.1 Calcium ovens	35
2.5.2 Magnetic field coils	37
3 Prototype surface-electrode ion trap	39
3.1 Trap design	40
3.1.1 Geometry	40
3.1.2 Operating parameters	42
3.2 Experimental apparatus	45
3.2.1 Trap packaging	45
3.2.2 Voltage supplies	46
3.3 Trapping results	46
3.3.1 Ion lifetime	47

3.3.2	Secular frequencies	48
3.4	Micromotion compensation	49
3.4.1	Out-of-plane micromotion detection	50
3.4.2	Theoretical modelling of modulated Raman effect	51
3.4.3	Experimental implementation	53
3.4.4	Micromotion compensation results	54
3.4.5	Comparison with other schemes	56
3.5	Heating rate measurement by Doppler recooling	56
3.5.1	Modified scheme for Lambda systems	56
4	Characterisation of a microfabricated trap	61
4.1	The traps	61
4.2	Voltage supplies	62
4.2.1	RF supply	62
4.2.2	DC supply	64
4.3	Operating parameters	65
4.4	Ion loading and lifetime	66
4.5	Micromotion compensation	67
4.6	Charging effects	68
4.7	Heating rates	70
4.8	Future improvements	71
4.8.1	Coated electrodes	71
4.8.2	Integrated filters	73
4.9	Comparison with other work	73
5	Pulsed laser cleaning	75
5.1	Anomalous heating in ion traps	75
5.1.1	Theoretical models	76
5.1.2	Cleaning methods	76
5.2	Experimental methods	78
5.2.1	Expected cleaning thresholds	78
5.2.2	Cleaning procedure	79
5.3	Results	80
5.3.1	Ion heating rate study	80
5.3.2	Electron microscopy study	83
5.4	Discussion	83
6	Qubits for microwave-driven quantum logic in $^{43}\text{Ca}^+$	89
6.1	Hyperfine qubits	89
6.2	The $^2S_{1/2}$ ground level of $^{43}\text{Ca}^+$	90
6.2.1	Structure	90
6.2.2	Transitions	91

6.2.3	Potential qubits	93
6.2.4	AC Zeeman shift	96
6.3	Doppler cooling	97
6.3.1	S-P-D level structure	99
6.4	Single-qubit operations	103
6.4.1	Qubit initialisation	103
6.4.2	Single-qubit gates	105
6.5	Qubit readout	105
6.6	Multi-qubit entangling operations	107
6.6.1	Driving motional sidebands with microwaves	107
6.6.2	The Mølmer-Sørensen gate	110
6.6.3	Gate speed	111
6.6.4	Alternative schemes	113
7	Design of trap and electronics for microwave-driven quantum logic	115
7.1	Trap design	115
7.1.1	Design goals	115
7.1.2	Trapping electrodes	116
7.1.3	Microwave conductors	118
7.1.4	Impedance matching and build-up resonator	120
7.1.5	Technical implementation	123
7.2	Microwave drive system design	129
7.2.1	Carrier drive	130
7.2.2	Initialisation drive	134
7.2.3	Sideband drive	134
7.2.4	Power amplification	135
7.2.5	RF/microwave diplexer	136
8	Construction and testing of microwave-driven quantum logic trap	141
8.1	Experimental apparatus	142
8.1.1	Vacuum system and trap electrical connections	142
8.1.2	Imaging system	147
8.1.3	Field coils	148
8.1.4	RF supply	148
8.1.5	DC supply	148
8.1.6	Optical pumping beams	149
8.2	Microwave testing	150
8.2.1	Trap testing	151
8.2.2	RF/microwave diplexer	154
8.2.3	Drive system	155
8.3	Testing with trapped ions	162
8.3.1	Ion loading and lifetime	162

8.3.2	Motional frequency stability	163
8.3.3	Micromotion	165
8.4	Microwave field measurements	165
8.4.1	Microwave field measurements by Rabi flopping	165
9	Conclusions	169
9.1	Summary	169
9.2	Towards microwave-driven entanglement	171
9.3	Future development: the 50-qubit processor	172
9.3.1	Trap	173
9.3.2	Laser interaction region	175
9.3.3	Single-qubit microwave operations	177
9.3.4	Crosstalk	178
9.3.5	Multiple qubit microwave operations	180
9.3.6	Scaling further	182
9.3.7	Conclusion	182
	Bibliography	185
A	Vacuum preparation	199
A.1	Vacuum compatible materials	199
A.1.1	Preparation	200
B	$^{43}\text{Ca}^+$ $S_{1/2}$ states and M1 matrix elements at intermediate field	201
B.1	$S_{1/2}$ hyperfine states	201
B.2	M1 transition frequencies and matrix elements	203
	List of abbreviations	207

1

Introduction

Raised on a diet of broken biscuits.

Pulp

Quantum computers allow information to be processed in a fundamentally different way to a classical computer because they rely on uniquely quantum-mechanical properties of superposition and entanglement. They will make certain classes of mathematical problems computationally tractable and they will allow for efficient simulation of quantum mechanical systems [NC00].

The necessary conditions for a quantum computer have been stated by DiVincenzo [DiV05] as:

1. A scalable physical system with well-characterized qubits
2. The ability to initialize the state of the qubits to a simple fiducial state, such as $|000\dots\rangle$
3. Long relevant decoherence times, much longer than the gate operation time
4. A “universal” set of quantum gates
5. A qubit-specific measurement capability

If quantum communication is also desired then a further two requirements can be added:

6. The ability to interconvert stationary and flying qubits
7. The ability faithfully to transmit flying qubits between specified locations

A wide variety of physical implementations for quantum computers have been proposed [LJL⁺10], including photons, neutral trapped atoms, quantum dots, solid state nuclear spins, liquid molecule NMR, superconducting circuits and trapped ions. Most of these implementations are experimentally at a very early stage with only one or two qubits having been demonstrated. The latter three implementations have been taken significantly further with simple, few qubit, algorithms having been demonstrated. Of these, liquid molecule NMR quantum computers were the earliest to be demonstrated but as yet there is no feasible proposal to scale them much beyond ten qubits. Superconducting circuits show great promise with entanglement of three qubits having recently been demonstrated [DRS⁺10]. Currently though, in terms of the scale of computations and gate fidelities achieved, trapped ions are the most mature implementation and the one which the research in this thesis is based on.

1.1 Ion trap quantum information processing

Quantum computing using trapped ions was first proposed in 1995 by Cirac and Zoller [CZ95]. The basis of their scheme is to use the collective motional modes of a string of trapped ions as a data bus to mediate quantum logic gates between ions. The entanglement of the internal state of an ion with an external motional degree of freedom was first demonstrated by Monroe et al. [MMK⁺95]. Since then there has been rapid growth in the field with over 30 groups currently working on trapped-ion QIP around the world. In this section I outline the basic

operations used in trapped-ion QIP, with reference to recent results which represent the state of the art. For a broader overview the reader is referred to recent review articles by Wineland [WL11] and Häffner [HRB08].

1.1.0.1 Ionic qubits

The qubits used in trapped-ion QIP come in two main types. In optical qubits the qubit transition is between a state in the ground level and a state in a metastable level (e.g. the 729 nm transition between states in the $S_{1/2}$ and $D_{5/2}$ levels of $^{40}\text{Ca}^+$). In ground level qubits the qubit transition is between two states in the ground level (e.g. the 3.2 GHz transition between an $F = 3$ and an $F = 4$ state in the $S_{1/2}$ ground level of $^{43}\text{Ca}^+$).

1.1.0.2 Initialisation and readout

Ions can be initialised to a well-defined internal state by optical pumping with circularly polarised light. This technique is well established with typical fidelities of 0.99 and is limited only by polarisation impurity of the pumping laser. This basic technique has been made more robust by employing a spectrally narrow transition to add frequency selectivity to the pumping and fidelities of > 0.999 have been achieved [RCK⁺06].

Readout in ions is achieved by driving fluorescence in such a way that only one qubit state scatters photons. This is done through the choice of laser polarisation and frequency [Lan06]. In ions with metastable levels one of the qubit states can be 'shelved' to one of these levels before driving fluorescence. Using this technique readout fidelities of 0.9999 have been achieved [MSW⁺08]. Another technique that can improve readout fidelity is quantum non-demolition (QND) readout. In this scheme the qubit state is transferred to an ancillary ion before being detected, a step which can be repeated many times to improve the fidelity. In this manner an Al^+ ion was read out using a Be^+ ion with a fidelity

of 0.9994 [HRW07].

1.1.0.3 Qubit coherence times

Typical ion qubits have coherence times of around a ms, with the main limitation being dephasing due to magnetic field noise. The use of π pulse sequences to dynamically decouple the qubit from the environment has been shown to increase coherence times by a factor of > 50 [Szw09]. The qubit can also be encoding in an entangled state of two ions which is immune to common-mode dephasing of the two ions, a so-called decoherence free subspace (DFS). These states have demonstrated to improve coherence times by a factor of over 100 and have been used in computations [MKV⁺09], though they increase the resources required in terms of qubits and gates. In ion species with nuclear spin there exist qubit transitions that are field-independent. In these qubits coherence times greater than 10 s have been measured [LOJ⁺05]. As gate times are typically in the 10-100 μ s range, decoherence in few-qubit ion experiments is rarely an issue.

1.1.0.4 Single-qubit gates

In both single and multi-qubit gates the qubit transition can be driven by lasers. In an optical qubit this can be done directly and in a ground level qubit using a two-photon Raman transition. The highest reported fidelity with Raman transition is 0.995 [KLR⁺08]. Pulsed lasers have also been used to drive Raman transitions, either in the form of frequency combs [HMM⁺10] or single pulses where single qubit gates can be driven in less than 50 ps [CMQ⁺10]. Ground state qubits can also be driven directly with rf or microwaves where much higher fidelities of 0.99998 have been reported [BWC⁺11].

1.1.0.5 Multi-qubit gates

The workhorse multi-qubit gates in ion traps are ‘warm-ion’ gates, such as the Mølmer-Sørensen gate, which do not require cooling of the ions to the motional ground state. The highest fidelity yet demonstrated for a two-qubit gate is 0.993 in a Mølmer-Sørensen gate at Innsbruck [BKRB08b]. This type of gate can entangle n ions in a single step, with a 14 ion GHZ state being the largest entangled state yet made [MSB⁺11]. Optical frequency combs have also been used to drive a Raman-transition entangling gate [HMM⁺10]. Recently microwave frequency magnetic field gradients have been used to drive a two-qubit gate of hyperfine ground level qubits directly [OWC⁺11]. This technique will form the basis of the later part of this thesis and is discussed there.

1.1.0.6 Addressing

Performing longer algorithms of many gates generally requires a method of individually addressing ions. As ions in the same trap are typically a few microns apart this can be done by tightly focussing a laser beam on the required ion [GRL⁺03]. Crosstalk with other ions is at the few percent level but this can be reduced with composite pulse sequences [HRB08]. Microwave or rf radiation has also been used with a strong magnetic field gradient applied across the ion string to give spectral addressability [WLG⁺09].

The addressing error can be eliminated completely by splitting the ions apart into separate traps hundreds of microns apart for addressing between multi-qubit operations. This splitting, shuttling and recombination will inevitably heat the ions but by using coolant ions of another species the ions can be recooled as required and computations of arbitrary length can be carried out (limited only by gate errors). At NIST this full set of methods has been used to implement two-qubit computation [HHJ⁺09].

1.1.0.7 Algorithms

Many quantum algorithms have now been demonstrated including, but not limited to, the Deutsch-Jozsa algorithm [GRL⁺03], the quantum Fourier transform [CBL⁺05], quantum teleportation [RHR⁺04] and Grover's search algorithm [BHL⁺05].

Another problem that trapped ions have been applied to is the simulation of other quantum systems. Magnetism in the quantum Ising model has been investigated directly in a nine ion system [IEK⁺11] and universal digital quantum simulation has been carried out using six qubits [LHN⁺11].

Basic error correction has been demonstrated using a three-qubit code to protect against one type of error [CLS⁺04]. Recently this has been extended to a measurement-free protocol that uses quantum rather than classical feedback to correct errors, again in a three-qubit code [SBM⁺11].

1.1.0.8 Photonic coupling

The entanglement of ionic qubits with photonic qubits offers a method of networking distant ion traps. In this way two ions in separate traps a meter apart have been entangled [MMO⁺07] in a probabilistic (but heralded) manner. The currently achieved entanglement rates are inconveniently slow ($\ll 1$ Hz) but large gains appear possible if fluorescence photons from the ion can be collected with high efficiency.

1.2 Focus of the thesis

Whilst impressive progress has been made in demonstrating all the required building blocks of an ion trap quantum computer there are two broad technical problems that impede the construction of a large scale device. Firstly the operations are not of a high enough fidelity that more than a few tens of operations

can be carried out sequentially. Secondly the current traps and control systems are not yet developed to the point where they could be scaled to hundreds of ions. Below I outline each of these areas in more detail. The various experiments described in this thesis attempt to make progress in both these areas.

1.2.1 Improving fidelity

If the fidelity could be improved (to $\gtrsim 0.9999$) then not only would longer computations be possible but quantum error correction could be employed to eliminate errors caused by gate infidelity entirely [Ste07].

1.2.1.1 Fundamental limitations

Some of the methods for qubit operations have no fundamental limits to their fidelity. An example of this is the technique of optical pumping to initialise ions in a particular ground state. Fidelity here is limited only by the polarisation purity available in the laser beam. Other methods are physically limited in the fidelity that can be obtained. An example of such a process is decay of the upper level of an optical qubit which is limited by the level's lifetime. In cases where the fundamental limit is not sufficiently high then new methods must be developed or a more favorable species of ion used (if the limitations are due to some property of the ion's atomic structure).

1.2.1.2 Trap limitations

The ion trap itself can limit the fidelity of operations in various ways that must be controlled through improved design, fabrication and preparation. The most important of these issues are outlined below and are the subject of investigation throughout this thesis. A textbook introduction in the context of ion trap chips for QIP is also available [ABLW11].

- Ion loss - The trapping lifetime is important as ion loss obviously means the loss of the qubit too. The ultimate limit to lifetime is background gas collisions but even in shallow traps this is typically many minutes and so the error introduced to a single operation is negligible. When transporting ions around multiplexed trap structures though, ion loss can be significant if the required electrode voltages are not repeatable (e.g. due to changing stray electric fields). Demonstrations of ion separation with no detectable failures [BCS⁺04] and 10^6 transports through a junction without loss [MHS⁺11] show that these issues are tractable with careful trap design.
- Uncompensatable micromotion - Ions are trapped at the pseudopotential minimum. This point is not necessarily a null of the rf trapping field though (eg. if there is an rf gradient in the axial direction). If this is the case the ion will have a driven motion at the rf frequency. This will limit how well the ion can be localised, will Doppler-shift lasers and make the ion susceptible to noise on the rf voltages.
- Stray fields - Charging of the trap and any other nearby structures will cause a shift of the ions' position and secular frequencies. The positional shift will cause micromotion by displacing the ion off the rf null, will reduce the repeatability of ion transport routines and will change the coupling to laser or rf/microwave fields that spatially vary. A shift in the secular frequency will introduce a detuning error on multi-qubit gates which rely on the stability of these mode frequencies. The stray fields can be compensated for through small adjustments to the electrode voltages but as this is time consuming only drifts on the timescale of minutes or longer can be conveniently compensated.
- Ion heating - During a multi-qubit gate the ions' internal states are entan-

gled with their motional state to some degree. Thermal heating of this motional state is therefore an important source of decoherence which causes gate errors. This heating is due to electrical field noise from the trap electrodes and is known as 'anomalous heating' as its source is not yet understood or well controlled. This noise scales unfavourably with distance d from the trap electrodes ($\sim 1/d^{-4}$) and becomes a significant source of gate error in traps with $d \lesssim 100 \mu\text{m}$.

- Nonlinearities - Anharmonic terms in the trapping potential cause amplitude-dependent shifts in the normal mode frequencies [HHJ⁺11]. We do not discuss these effects in this thesis as the coherence of the motional states is limited to a greater extent by the ion heating, though they will become increasingly important as traps get smaller and higher fidelities are required.

1.2.1.3 Control limitations

The leading causes of infidelity in the experiments thus far realised are not fundamental or due to the traps used but due to the technical limitations of the control schemes utilised. Some of these limitations, such as producing less noisy voltage supplies are clearly solvable with currently available technology. That they have not yet been satisfactorily solved is due only to the lack of required electronics expertise and resources within the small groups of specialised atomic physicists that typically run ion trap experiments. Other limitations such as noise on ultrastable lasers represent areas where technical improvements are not readily available. In these cases it is likely that qubit operations that are intrinsically more robust to these errors need to be developed to reduce the magnitude of technical improvement required. A representative selection of typical issues is outlined below.

- Trap voltages - Electrical noise on electrode voltages at the ions' secular

frequencies (or motional sidebands of the rf trapping frequency) will heat the ions. Amplitude stability of voltages is required to keep ion positions and motional frequencies stable.

- Magnetic field stability - Fast noise causes qubit dephasing and slow drifts cause detuning errors on qubit operations.
- Lasers, microwaves and rf - Amplitude, phase and frequency noise and drifts all contribute to infidelity. Additionally for lasers, amplified spontaneous emission and beam pointing errors are also an issue.
- Photon collection efficiency - A lower collection efficiency decreases the fidelity of qubit readout operations.

1.2.2 Improving scalability

The first requirement of a large scale QIP device is to produce an array of traps that can handle large numbers of ions. With the advent of surface electrode traps that can be microfabricated this requirement has largely been met and traps, with over a hundred zones have been tested [AUW⁺10], though much optimisation remains to be done.

The next main remaining challenge is to scale the control systems. Scaling the classical computer control and trap drive electronics to hundreds of ions should not pose any particular issues. The laser control however poses a much greater challenge. Firstly the laser sources need reducing in size and complexity. Secondly a way of getting many laser beams into the trap array needs to be devised as the current system of bulk optics around the vacuum can is clearly not scalable. Both of these issues will be discussed in the conclusion.

1.3 Structure of the thesis

Chapter 2 describes the experimental apparatus required to run an ion trap experiment. The apparatus includes lasers, optics, vacuum systems, trap fabrication, trap drive electronics and the control systems to integrate them.

Chapter 3 describes the design and operation of a surface electrode trap. A modified method of micromotion compensation that is suitable for use in surface traps is demonstrated. We also demonstrate a repumping scheme that simplifies heating rate measurements in 3-level systems.

Chapter 4 describes the testing of a trap designed in collaboration with Sandia National Laboratories and fabricated by them using semiconductor microfabrication techniques. We investigate in detail electrode charging, heating rates and micromotion compensation.

Chapter 5 describes a method of in-situ pulsed laser cleaning of ion traps to reduce the ion heating rate. This is demonstrated on one of the traps described in chapter 4.

Chapter 6 gives an overview of hyperfine qubits in the ground state of $^{43}\text{Ca}^+$ and evaluates their suitability for a microwave-driven multi-qubit scheme.

Chapter 7 describes the design of a trap incorporating microwave circuitry to drive quantum logic gates. The design of the associated microwave drive system is also described.

Chapter 8 describes construction of the experiment designed in chapter 7 and experiments with trapped ions to characterise the trap as well as electrical testing of the microwave systems.

Chapter 9 summarises the thesis and outlines the experimental steps required to complete the experiment begun in chapter 8. Finally we describe future prospects for scaling the methods described in this thesis to a larger quantum information processor.

2

Experimental apparatus

You can arm me to the teeth
You can't make me go to war

The Mountain Goats

This chapter details the experimental apparatus and experimental techniques that are common to the three sets of experiments that are described in subsequent chapters. Where a particular piece of apparatus is used only on a single experiment we detail it in the relevant chapter. Figure 2.1 gives a block-diagram overview of the apparatus.

2.1 Laser systems

For the experiments described in this thesis we use laser sources to drive all five of the electric dipole transitions between the $4S$, $4P$ and $3D$ levels in Ca^+ (see fig 2.2b). We also require two further lasers to photoionize neutral Ca in order to load the trap (see fig 2.2a). Calcium is a particularly convenient ion with regards to lasers in three respects. Firstly, simple diode lasers are available at all the required wavelengths. Secondly, optical fibres with low loss (over tens of meters) are available at these wavelengths. Finally, all the transitions lie in two narrow wavelength ranges (389-423 nm and 850-866 nm) which allows standardisation

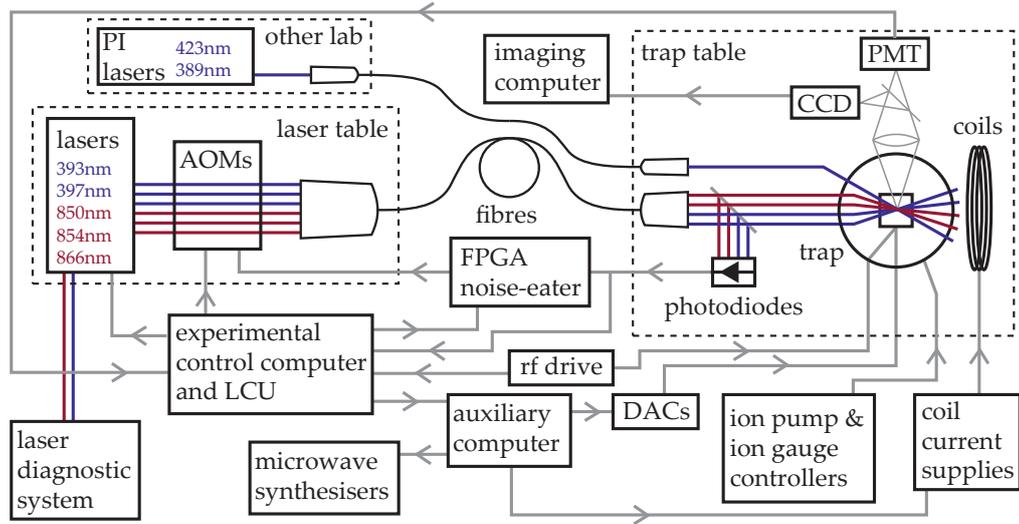


Figure 2.1: An overview of the experiment. Each section of the diagram is covered in more detail in this chapter.

of much of the optics. These two wavelength ranges will henceforth be referred to as ‘UV’ and ‘IR’ for convenience.

2.1.1 Photoionization lasers

The two-step resonant photoionization technique we use is described in previous work by the group [LRH⁺04]. The lasers are located in another lab where they are shared with other experiments. The 423 nm laser is a Toptica DL100 extended cavity diode laser and the 389 nm laser is a free-running diode on a home-made mount. Both are combined, with the same linear polarization, on a dichroic mirror (Semrock FF409-Di02) and coupled into a single-mode fibre. The lasers are switched with a shutter before the fibre.

2.1.2 Ca⁺ lasers

All five of the Ca⁺ lasers are extended cavity diode lasers in a Littrow configuration. The lasers and their associated temperature, current and piezo control electronics are commercial systems from Toptica Photonics AG (DL100). Table 2.1

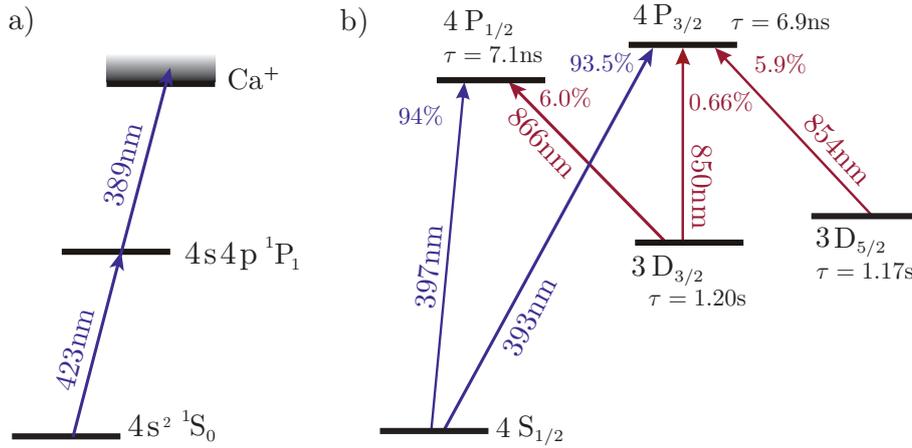


Figure 2.2: a) Transitions used to photoionize neutral calcium. b) Low-lying electric dipole transition wavelengths, excited state lifetimes and branching ratios in Ca^+ .

Table 2.1: The maximum power of each of our laser diodes after its feedback grating (bare diode for 389 nm).

Laser λ (nm)	Power (mW)
388.9	5
422.79	16
393.48	3.7
396.96	16
850.04	178
854.44	45
866.45	41

summarises the laser diode output powers and figure 2.3 shows the optical configuration of each system. Optical isolators (> 30 dB isolation) are placed immediately after each laser to prevent back-reflections from components further down the beam path disturbing the laser by optical feedback. The IR lasers are locked to the side of a reference cavity fringe using the PID controllers incorporated into their control electronics, which feed back on both the current to the diode and the piezo on the laser's external cavity grating. The UV lasers are locked to reference cavities using Pound-Drever-Hall (PDH) locking [DHK⁺83]. The modulation for the lock comes from a resonant EOM running at 81 MHz and the PDH error signal is generated using electronics developed by the group

2. EXPERIMENTAL APPARATUS

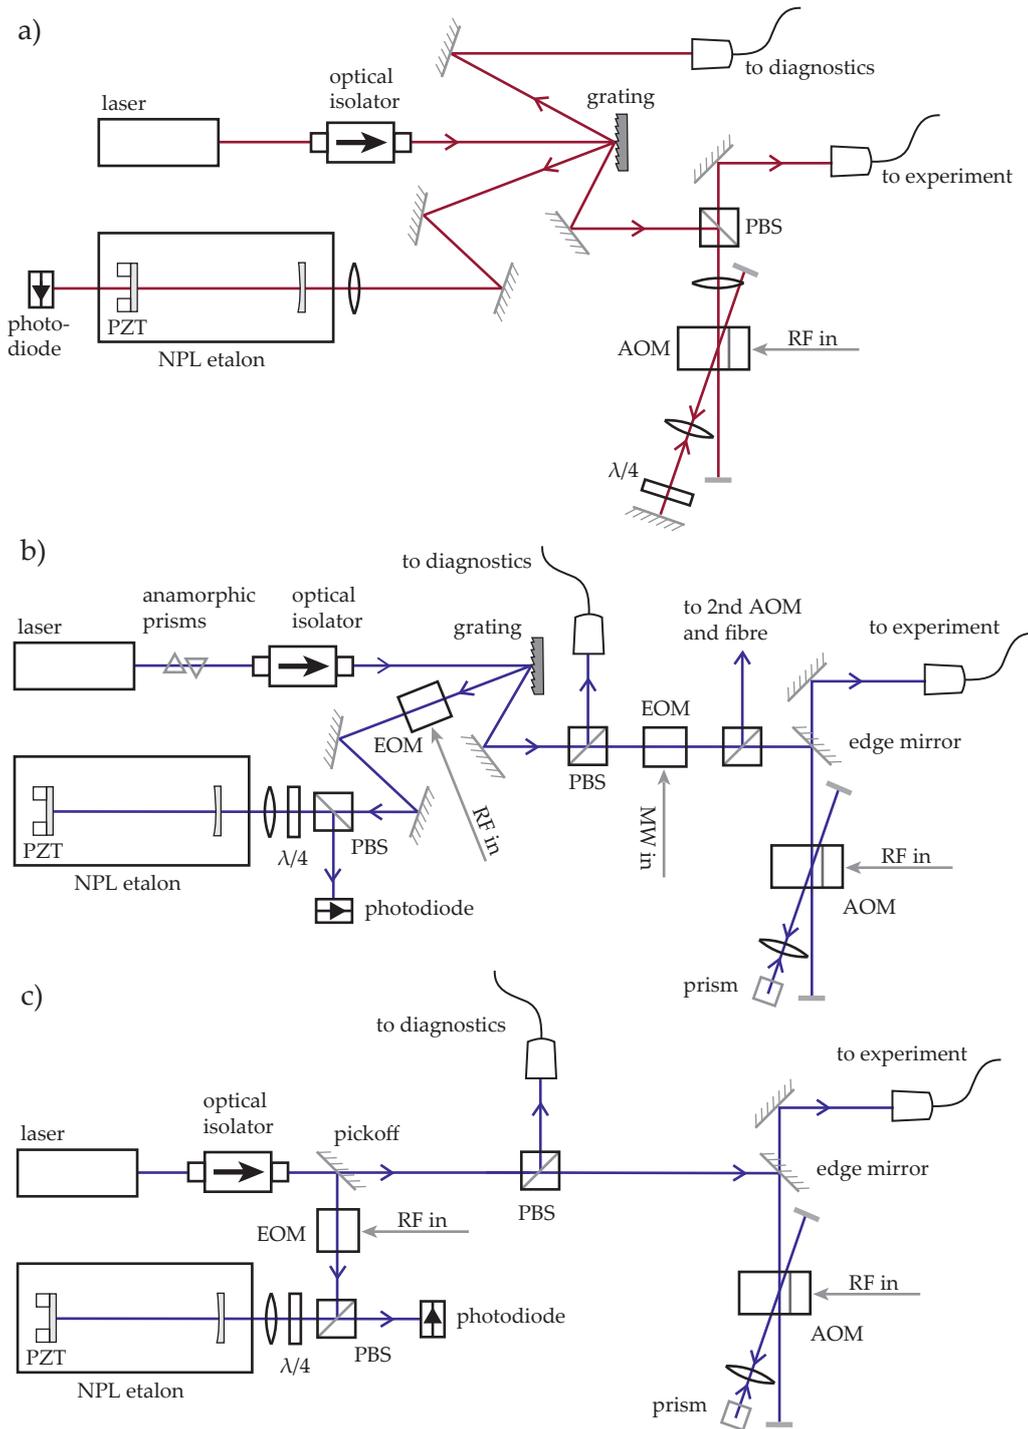


Figure 2.3: Optical layout for a) IR lasers, b) the 397 nm laser and c) the 393 nm laser. The 397 nm beam is split at the point indicated with a fraction of the power going to a second AOM and fibre to provide a separately switchable optical pumping beam. The schematic has been simplified for clarity by omission of lenses, half-wave plates and some mirrors. Figure modified from [Cur10].

[Kei07]. This error signal is then fed into the laser's PID controller. The low-drift (<500 kHz/hr) reference cavities are commercially produced by NPL and incorporate a piezo which allows the frequency of the laser locked to them to be scanned by the experimental control computer (see section 2.2.1). A blazed diffraction grating is included in the the beam path of most of the lasers to prevent amplified spontaneous emission (ASE) at other wavelengths reaching the ions. Without the gratings, ASE from the 397 nm or 866 nm lasers near 393 nm and 850 nm could shelve the ion to the $D_{5/2}$ level and ASE from the 850 nm or 866 nm laser near 854 nm could deshelve an ion shelved in the $D_{5/2}$ level. An EOM¹ in the 397 nm beam path is used to add sidebands to the laser which span the 3.2 GHz hyperfine splitting in this transition when working with $^{43}\text{Ca}^+$.

2.1.2.1 Acousto-optic modulators

Acousto-optic modulators are used to switch and to control the amplitude of all the Ca^+ lasers. The IR AOMs² run at 85 MHz and are specified at 80% first-order diffraction efficiency per pass. The UV lasers use either 85 MHz³ or 200 MHz⁴ AOMs depending on the experiment. The 85 MHz ones have a better specified diffraction efficiency (80% into first-order as opposed to 70%) but any unwanted scatter off the AOM crystal that reaches the ion will be closer to resonance. To obtain better extinction when the AOMs are off we use them in a double-pass configuration (see figure 2.3). As the IR AOMs are polarization insensitive this is achieved by rotating the polarization for the second pass. The UV AOMs require a certain polarization for maximum efficiency so a prism and lens arrangement is used to displace vertically the second pass which is then picked off with an edge mirror. The AOM drivers include a TTL digital switch and DC-coupled AM and FM inputs. AOM extinction is limited by rf leakage from the driver but

¹New Focus 4431M resonant EOM, tuneable from ≈ 2.4 to ≈ 3.7 GHz

²IntraAction Corp. ATM-851A2 with DE-852EM26A driver

³IntraAction Corp. ASM-851.5B8 with DE-853EM26A driver

⁴IntraAction Corp. ASM-2001.5B8 with DE-2003EM26A driver

can be improved if required by applying a digital pulse to the FM input [Szw09]. This both detunes the laser from the atomic transition and reduces the leakage by moving the driver frequency outside of the AOM's bandwidth.

2.1.2.2 Fibres

The lasers are located on a dedicated optical table and single-mode polarisation-maintaining fibres carry the light to the ion trap optical table. This table has two separate ion trap experiments on it. The experiments cannot be run at the same time as there is only one experimental control system and one set of lasers but switching between them is trivial. For the IR light we use PMC-780⁵ fibres and for UV we use PMC-S405Si⁵ or PMC-360Si⁵ fibres. Angle-polished FC connectors (FC-APC) are used to prevent problems with optical feedback which, even with an optical isolator immediately after the laser, was found to occur with non-angle-polished fibres. We often couple more than one wavelength into the same fibre for convenience. When two wavelengths with orthogonal polarizations are required this can be achieved without loss by combining beams on a polarising beamsplitter (PBS) before the fibre. 393 nm and 397 nm beams with collinear polarization are required for the UV σ beam path (see section 8.1.6) here an interference filter⁶ at 1.5° incidence is used as a dichroic mirror to combine the beams with <5% loss in each beam. For the IR lasers we use non-polarizing beamsplitters (NPBS) when collinear polarization is required. This is less efficient but we have a surplus of power in the IR (see table 2.1).

2.1.3 Laser diagnostic system

Figure 2.4 shows the laser diagnostic system. The LabView display for this system (see figure 2.5) indicates when the laser is at the required wavelength, warns if a laser jumps to another wavelength and, from the optical spectrum analyser

⁵Schäfter + Kirchoff GmbH

⁶Semrock FF01-387/11

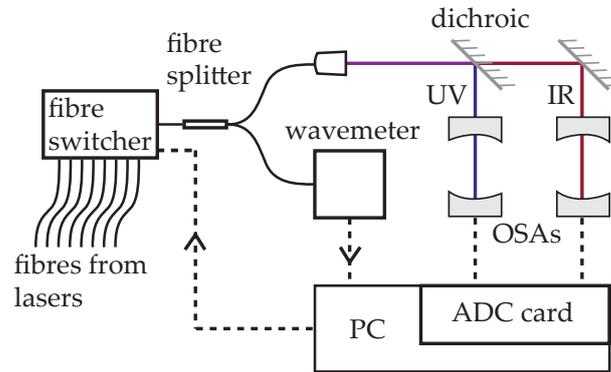


Figure 2.4: Laser diagnostic system

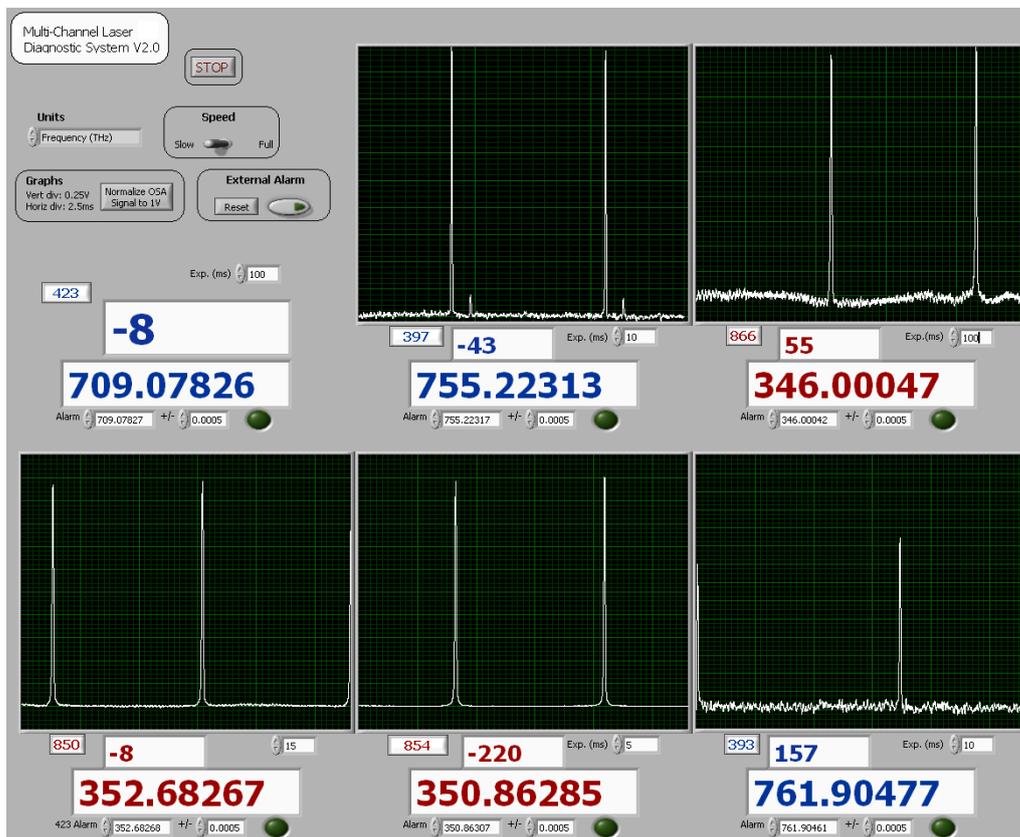


Figure 2.5: Screenshot of the laser diagnostic system running. The update rate is a few Hz divided by the number of channels in use. The lower number for each laser is the absolute frequency in THz and the upper number is the detuning from the transition in MHz. Multimode lasing on the 397 nm laser is clearly visible, indicating that the current requires adjustment.

(OSA) trace, indicates if a diode is lasing on more than one longitudinal mode. The wavemeter ⁷ is a Fizeau interferometer type and the OSAs ⁸ are scanning confocal cavity types. The signal output from their control electronics is acquired with an ADC card ⁹. The diagnostic light ($\gtrsim 200 \mu\text{W}$ is required) is carried from each laser to the fibre switcher ¹⁰ using $50 \mu\text{m}$ -core multimode patch cables. The fibre splitter¹¹ is a 90:10 splitter made of the same fibre with the 10% portion going to the wavemeter. For more information see [All07].

2.1.4 Trap optics

Figure 2.6 shows the optical arrangement for delivering the cooling laser beams from the fibres to the ion. We combine all the main Doppler cooling lasers into a single beam path to simplify alignment. Also shown is the 45° IR beam described in section 3.4. Experiments in chapter 8 make use of additional beam paths for qubit initialisation and readout and are described in section 8.1.6. The photodiodes¹² are for the noise eater (see section 2.2.3) and also allow the laser powers to be logged. The photodiodes have a short length of tubing attached to the front which contains a short focal length lens to reduce the spot size and an interference filter¹³ to block ambient light.

2.1.4.1 Collimation and focussing

In the experiments described in chapter 3 the ion was $150 \mu\text{m}$ from the trap surface. Initially a collimator with a moulded asphere ¹⁴ was used for the both the 397 nm and photoionisation fibre outputs. In later traps though the ion was

⁷High-Finesse WS-7 with 10 MHz resolution and 60 MHz absolute accuracy

⁸The UV OSA is a Toptica Photonics FPI100 with a FSR of 1 GHz and a finesse of ~ 100 . The IR OSA is a Technical Optics SA-7.5 with a FSR of 2 GHz and a finesse of ~ 50

⁹National Instruments PCI-2515 with eight 16-bit, 1.25 Msample/s ADCs

¹⁰High Finesse MC-8

¹¹Thorlabs FCMM50-90A-FC

¹²Thorlabs PDA-36A

¹³Thorlabs FB390-10, FB400-10, FB850-10 or FB870-10 depending on wavelength, all 10 nm FWHM

¹⁴Thorlabs PAF-2-X-A, $f=2 \text{ mm}$

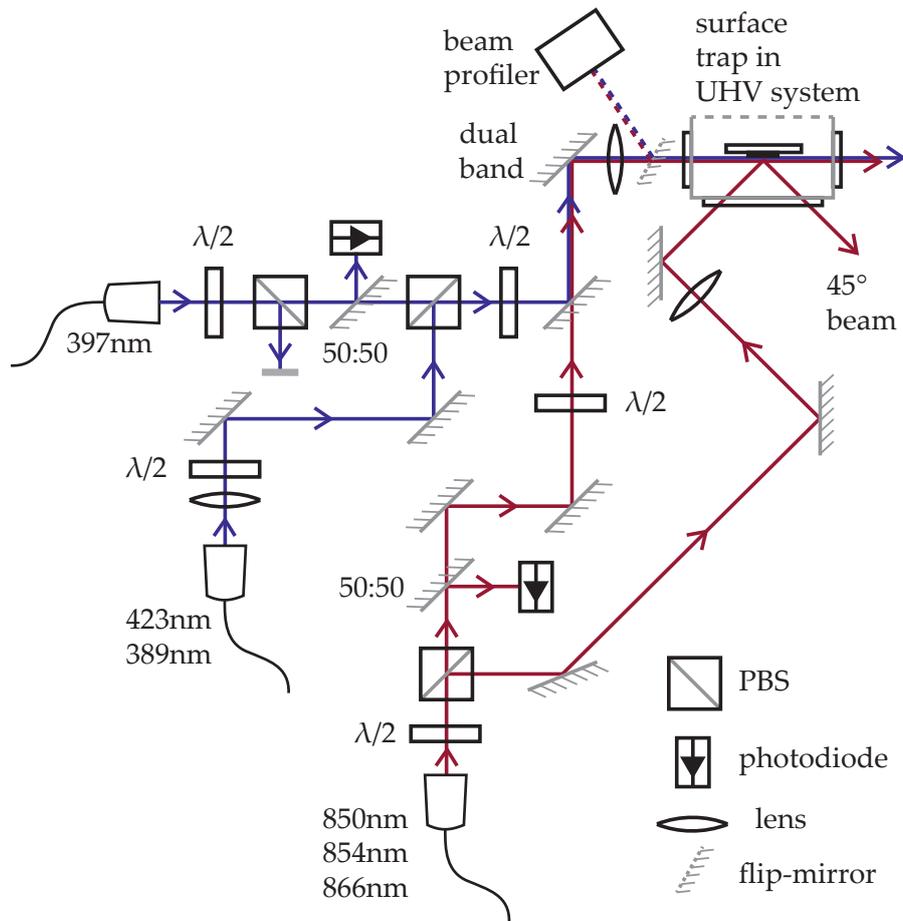


Figure 2.6: Optical arrangement for the main Doppler cooling and photoionization laser beams used on the trap table (with slight alterations) for all the experiments described in this thesis.

75-84 μm from the trap surface which required higher quality optics to reduce 397 nm scatter. Firstly a multi-element $f=14.5$ mm collimation lens¹⁵ was used which gives no noticeable aberration, in contrast to several collimators with single aspheric lens that we tried which all produced visible halos around the beam. For the PI fibre we retained the aspheric collimator and to match the spot size of the 397 nm beam we set the asphere to give a diverging beam and then collimated it using a 300 mm achromat about 250 mm from the fibre. As IR scatter is not critical we simply used a 30mm achromatic doublet for collimation. As this

¹⁵CVI Melles Griot GLC-14.5-8.0-405, a diffraction-limited meniscus and achromatic doublet combination with NA=0.28

2. EXPERIMENTAL APPARATUS

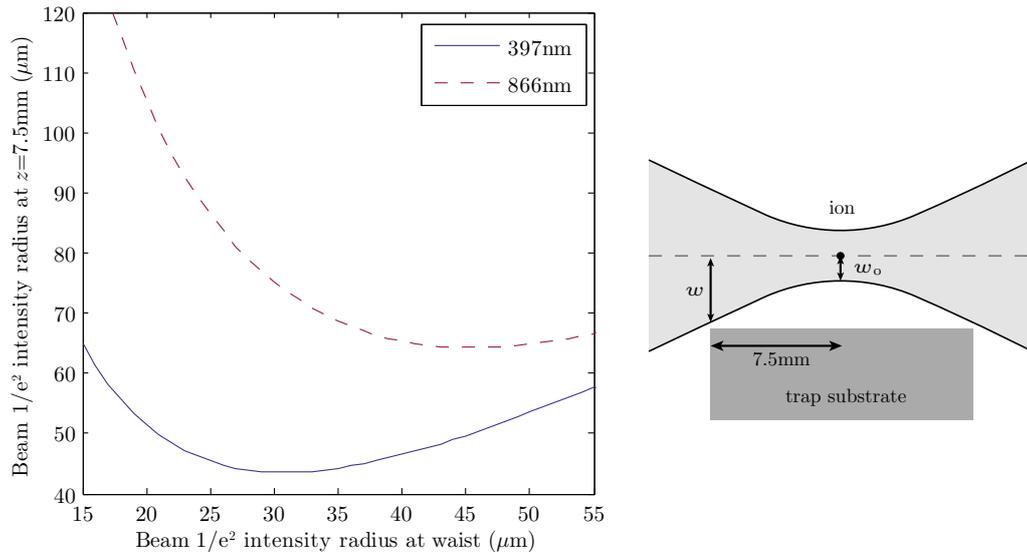


Figure 2.7: Radius ($1/e^2$ intensity) of the 397 nm and 866 nm laser beams at the edge of the trap substrate ($z=7.5\text{ mm}$) for a given beam waist at the ion's position ($z = 0$).

has twice the focal length of the 397 nm collimator it gives a similar sized spot at the ion. We have found in the past that glass optics can contribute significantly to background scatter at 397 nm so fused silica optics were used where possible. Focussing onto the ion was achieved with a 200 mm lens. This gives a beam waist at the ion of $\approx 30\ \mu\text{m}$, which minimises the beam size at the edge of the substrate (see figure 2.7) in an effort to minimize the scattered light background. In chapter 3 this was a simple plano-convex lens but from chapter 4 onwards achromatic doublets were used. In all the traps in this thesis the signal to background ratio from scattered light for a cold ion of better than 100 was readily achievable. With the improved optics in chapters 4 and 8 a signal to background of better than 200 was achieved despite the ion being only 75 – 85 μm from the electrodes.

2.1.4.2 Alignment onto ions

In order to align and to focus the main cooling beams onto the ion a scanning slit beam profiler¹⁶ is used. The beams are diverted to the profiler using an aluminium mirror on a flip-mount and care is taken to place the beam profiler at the same distance from the final lens as the ion. First the 397 nm laser is turned on and the lens position adjusted to minimise the spot size on the profiler (i.e. to put the waist at the ion). Next the beam is aligned in the direction parallel to the trap by grazing it across the trap surface so that it is easily visible on the imaging camera. The trap structure can then be used as a reference to direct the beam through the trapping location with high accuracy. Finally the beam has to be aligned out of the plane by first making it parallel to the trap and then translating it away from the trap by the ion-trap distance. This is done by moving the beam towards the trap using the final steering mirror until the first sign of clipping is apparent on the beam exiting the vacuum system. The flip mirror is then inserted and the beam position on the profiler recorded. The mirror before the final mirror is then adjusted to alter the tilt of the beam. The first step is then repeated to see if the beam gets closer to the trap before it clips. A minimum will exist which corresponds to the beam being parallel to the trap. The beam is then moved just off the trap surface. In chapter 3 the final mirror is silvered. However, after this we moved to a dual-band dielectric mirror¹⁷ which gave less loss and scatter. We also put this mirror in a gimbaled mount with micrometer adjusters¹⁸ which made beam alignment relative to the trap much easier. The other beams are overlapped with the 397 nm beam at the ion using the beam profiler and checking by eye that they are overlapped further down the beampath. The collimators for the other lasers are also adjusted to ensure they come to a focus at same place as the 397 nm beam.

¹⁶Thorlabs BP104-UV, 0.5 μm resolution

¹⁷Thorlabs UFM10R

¹⁸Thorlabs GM100

2.2 Experimental control

The experiment is controlled by two computers and an FPGA. The first computer is the experimental control computer which runs all the time-critical aspects of the experiment using specialised hardware including the laser control unit (LCU). The second, auxiliary, computer is a more flexible PC which can run a variety of software and hardware as the experiment dictates. For the work in chapter 4 onwards the FPGA was added to the auxiliary computer for fibre noise-eating. In addition there are two further computers on the experiment for diagnostic purposes. One runs the laser diagnostic system described above and one runs the CCD camera that images the ion.

2.2.1 Experimental control computer

The main experimental control program is written in Turbo Pascal and runs on the experimental control computer under MS-DOS. The program utilises attached hardware (see figure 2.8) to carry out four main functions. These are switching laser and microwave sources, scanning laser frequencies, counting photons and logging experimental parameters. The switching is done through TTL outputs on a PCI card ¹⁹ which are connected to the AOM controllers (to switch beams) or microwave switches. These outputs can be controlled with a 0.1 μ s resolution using the laser control unit (LCU), a system of hardware timing utilising the counters on the card which is described in detail in [Kei07]. Laser frequencies are scanned using a DAC card ²⁰ which is connected, via high voltage drivers, to the reference cavity piezos. Experiments can be synchronised with the 50 Hz mains supply using the line trigger which produces a TTL pulse at the same point in every mains cycle. The experimental control computer logs laser powers and other relevant signals using its ADCs.

¹⁹Measurement Computing PCI-DAS1200 PCI card with sixteen 12-bit ADCs, two 12-bit DACs, 24 TTL digital input/outputs and six 16-bit counters

²⁰Advantech PCL-727 ISA card with twelve 12-bit DACs

2.2.1.1 Photon counting

Photons from the ion are counted using an Electron Tubes Ltd P25PC-12 PMT. Figure 2.8 shows the experimental control computer's counter hardware²¹ used to count the photons. Their arrival time can also be recorded with respect to a fixed point in the trap rf cycle, information which is necessary for implementing the micromotion detection scheme described in section 3.4. This is done using a time-to-amplitude convertor²² (TAC) which outputs a voltage proportional to the time between registering a photon and a fixed point on the rf cycle. This voltage is then read into the experimental control computer using an ADC. The reference point on the rf cycle is produced by a constant fraction discriminator²³ connected to the trap rf synthesizer. For the 10 MHz photon counters to work correctly the 8 ns TTL pulses from the PMT must first be increased in duration to 50 ns using a 'pulse stretcher'²⁴. This pulse stretcher module also incorporates a 1.6 μ s delay generator between the TAC's 'valid conversion' output and the ADC pacer that is necessary to trigger the ADC at the correct time for sampling the TAC output pulse.. For further details see [Mye07].

2.2.2 Auxiliary computer

Initially for the work in chapters 3-5 the auxiliary computer only controlled the DACs which provided the trap dc electrode voltages. These were controlled from a LabView interface. For the work described in chapter 8 we added a GPIB interface²⁵ in order to interface with the microwave synthesisers and magnetic field coil power supplies. We also added a serial (RS-232) communication link between this computer and the experimental control computer. This allows the experimental control program to request changes in experimental parameters

²¹Advantech PCL-836 ISA card with six 16-bit counters

²²ORTEC model 566

²³ORTEC model 584

²⁴Oxford Physics Central Electronics Group

²⁵Prologix GPIB-USB

2.2.3 Fibre noise-eating FPGA system

Fibres add intensity noise in two ways. Firstly, any small variations in beam pointing on the fibre input will cause large changes in the amount of light coupled into the fibre. Secondly, changes in temperature and stresses on the fibre alter its birefringence which causes the polarization of the output beam to change (use of PM-fibre reduces, but does not eliminate this effect). This polarisation noise is converted into amplitude noise by placing a polariser after the fibre as we require well-defined polarizations for our experiments. The characteristic timescales of this noise varies from hours (e.g. day-night thermal drift) to kHz (e.g. mechanical resonance of optical components) and amplitudes of order 10% are typical. This fibre noise can be removed by measuring the power after the post-fibre polariser and using a feedback circuit to modulate the rf power to the pre-fibre AOM (the AOM drivers have an AM input for this purpose). If the laser were to be on constantly, a simple analogue PID feedback loop would suffice; however most experiments require the lasers to be operated in a pulsed manner. In an analogue PID loop the integral part of the circuit would rail as soon as the laser was switched off. Then, when the laser was switched back on the AOM power would be set to the wrong level which would cause a transient in the power. Therefore a way of “pausing” the PID loop is required when the laser is off. There is also a further complication that some pulses are so short (\sim few μ s) that a feedback loop would have to be very fast to operate during a single pulse. Alternatively, provided we only want to remove slower noise, the power could be sampled once per pulse or sampled between experimental runs. Analogue circuitry that could meet these requirements would be complicated and hard to reconfigure so a digital feedback circuit was chosen instead. In order to obtain fast, deterministic execution and easy reconfigurability the loops are implemented on an FPGA. For convenience we use a commercial PCIexpress

card²⁶ which has 8 ADCs and 8 DACs interfaced to an FPGA. This card can be programmed in LabView which avoids the need to learn a specialist hardware description language (HDL). Figure 2.9 shows the system. A breakout box was made to interface the FPGA card to the experiment. The breakout box is built into a 19" rack mount case and connects to the FPGA card through a screened cable²⁷. $50\ \Omega$ output buffers²⁸ and input termination must be provided as both the AOM driver and photodiode²⁹ are $50\ \Omega$ devices connected via $50\ \Omega$ BNC cables. A lowpass filter³⁰ is also included to remove any high frequency noise (from the PC power supply or DAC digitisation noise) above the feedback loop's bandwidth (up to 100 kHz). The system has been shown to reduce the amplitude of laser intensity noise to $<1\%$ which is negligible for the experiments performed in this thesis. The noise suppression at 2 kHz was measured to be -20 dB. For more details see [Rob10].

2.3 Imaging system

Typically fluorescence photons from the ion are collected using a large numerical aperture objective lens and counted by the PMT. Figure 2.10 shows the system used in chapters 3-5, which was mounted horizontally, along the optical table. The main objective is a Nikon ED PLAN 1.5X wide aperture compound lens with N.A. ≈ 0.29 , which corresponds to 2.1% of the full 4π solid angle. This lens images the ions with a magnification of 8x onto a pair of adjustable slits. These slits can be opened to allow imaging of the entire trap structure using the camera, useful when aligning lasers and the imaging system, or closed down to a small rectangle around the ion to reduce background scatter on the PMT. This

²⁶National Instruments PCIe-7852R with Virtex-5 LX50 FPGA clocked at 40 MHz, eight 750 kHz, 16-bit ADCs and eight 750 kHz, 16-bit DACs

²⁷National Instruments SHC68-68-RMIO

²⁸Intersil HA5033 unity gain video buffer

²⁹Thorlabs PDA-36A

³⁰Maxim MAX275 active 4th-order filter with centre frequency of up to 300 kHz

shutter is then re-imaged with 1x magnification onto the PMT or, by rotating in a beamsplitter, onto the CCD camera. The system was built by Alice Myerson and a precise calibration of the optics can be found in [Mye07]. For the work in chapter 8 a new, vertically mounted imaging system was built which is described in that chapter.

The net detection efficiency η of the imaging system (including the quantum efficiency of the PMT) was measured to be 0.24% using a pulsed method [SKDB10]: the ion was shelved in $D_{3/2}$ by an optical pumping pulse from the 397 nm laser alone, then deshelled by the 866 nm laser alone to emit a single 397 nm photon which is detected with probability η . This sequence is repeated many times to measure η . The detection window must be sufficiently long to ensure precession out of dark states in $D_{3/2}$ (which are formed for any static polarization of the 866 nm light [BBHK00]); the precession timescale is $\sim 2 \mu\text{s}$ for our conditions.

2.3.1 Filters

During normal operation two interference filters³¹ are used in front of the PMT to reduce background counts from room lights and the UV lasers other than the 397 nm cooling laser (the PMT is insensitive to the IR lasers). When observing neutral Ca fluorescence in order to check for correct oven operation and 423 nm laser wavelength and alignment, we use a different filter³² to reduce background counts from room lights. The group has also used another set of filters³³ to block 397 nm light whilst observing fluorescence at 393 nm. For the camera a simple coloured glass³⁴ filter to block the IR light suffices. This is removed to align the 45° IR beam by observation of the beam striking the trap surface.

³¹Semrock FF01-377/50-25 and FF01-406/15-25 with combined 397 nm transmission of 86%

³²Semrock FF01-435/40-25, with 423 nm transmission of 93%

³³Two Semrock FF01-387/11-25 filters with net transmission of 78% at 393 nm and 0.3% at 397 nm

³⁴Schott BG38

2.3.2 Imaging computer

The imaging computer displays real-time images of the ions in the trap when the camera beamsplitter is in place (see section 2.3). This is a useful diagnostic as it shows if the ions are crystallised (and thus cold) and, if they are, the number of ions in the trap. For the work in chapter 3 we used an Andor DV437³⁵ CCD camera with thermo-electric cooling. From chapter 4 onwards we switched to a newer Andor Luca DL-604³⁶ electron multiplying (EMCCD), which offers faster readout.

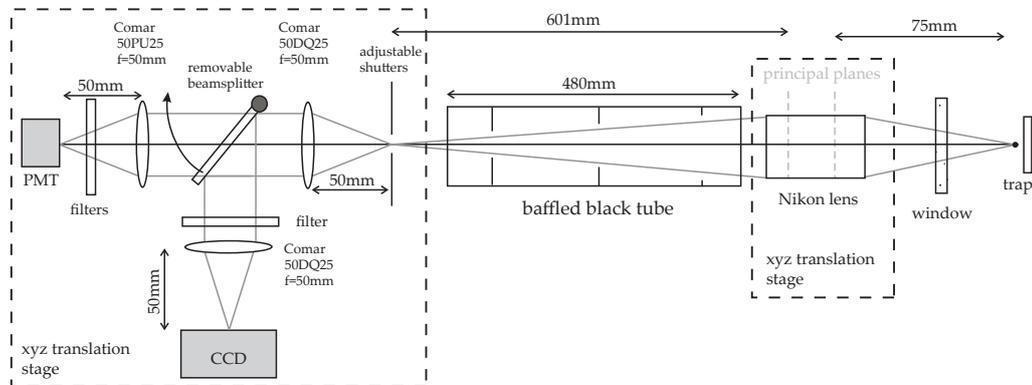


Figure 2.10: The horizontal imaging system. Figure modified from [Mye07].

2.4 Trap fabrication

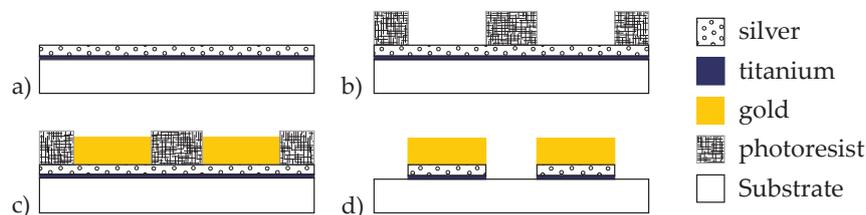


Figure 2.11: Fabrication stages (see text).

³⁵515×515 pixels, 13×13 μm pixel size, 1 MHz readout, 60% QE, PCI interface

³⁶1004×1002 pixels, 8×8 μm pixel size, 13.5 MHz readout, 40% QE, USB 2.0 interface

The traps used in chapters 3 and 8 were fabricated using a process we developed based on that of I. Chuang's group at MIT [Lab08] but using the electroplating process developed in [KMK⁺04]. The fabrication was carried out in the physics department's class 10,000 cleanrooms. We fabricated the traps on polished single-crystal substrates of quartz or sapphire. The substrates were first thoroughly cleaned in piranha etch (10 parts H₂SO₄, 1 part H₂O₂) at 100°C for 3 minutes then rinsed in high quality de-ionized (DI) water³⁷. The substrates were then glow-discharge cleaned and a 15 nm titanium layer was evaporatively deposited to ensure good adhesion before evaporating on a 200 nm silver seed layer (figure 2.11a). The substrates were then cleaned in acetone, isopropyl alcohol (IPA) and DI water followed by a bake at 115°C to dry. To promote good photoresist adhesion hexamethyldisilazane (HMDS) primer was used. Initially we diluted the primer 1:4 with propylene glycol monomethyl ether acetate and puddled it on the substrates for 10s before being spun off at 5000 rpm for 45 seconds. We switched to vapour priming the substrate on a hotplate at 115°C though for later work as it reduced the chance of leaving residue or particulates on the trap. The photoresist³⁸ was puddled onto the substrate and spun off at either 4000 or 5000 rpm for 45 seconds to give a nominal thickness of 6.5 or 5 μm. After being spun, the photoresist was baked for 2 minutes at 115°C before resting at least 30 minutes to rehydrate the photoresist. A mask aligner was used with a chrome on glass photomask³⁹ for the photolithography of the trap structure. This was done in two stages. The first exposed and removed the edge-bead caused by photoresist surface tension to ensure good mask contact and to prevent undesirable regions of bare substrate on the trap. The second defined the trap structure (figure 2.11b). After each stage there was a 2 minute development (using Microposit MF CD-26) and a DI water rinse.

³⁷Fisher Scientific LC-MS grade, <0.0001% residue

³⁸Rohm and Haas Megaposit SPR 220-7.0 positive resist

³⁹JD Photo-Tools, 128,000 dpi resolution raster imaged mask

A fully bright sulphite gold electroplating process (Metalor ECF 63) plated the gold electrodes up within the photoresist mask (figure 2.11c). Plating was carried out using 200 ml of solution in a beaker on a stirring hotplate. The trap and anode are suspended in the beaker facing each other at a distance of a few cm. The anode is pure platinum foil with an area of $\sim 500 \text{ mm}^2$. Plating was carried out at 55°C in a constantly stirred solution at a rate of $\sim 1 \mu\text{m}$ of gold every 7.9 mins. After electroplating the photoresist was removed using acetone, IPA and DI water. Then the silver seed layer was etched using $\text{NH}_4\text{OH}:\text{H}_2\text{O}_2:\text{H}_2\text{O}$ (1:1:4) for 15 seconds. The final trap used in chapter 3 was inspected and the gold thickness measured using a scanning electron microscope (SEM, see figure 2.12) to be $\sim 2.7 \mu\text{m}$. The surface quality obtained was mirror-like and no surface features could be found by the SEM which indicates an upper bound on surface roughness of around 50 nm, in agreement with atomic force microscope studies of similar devices [KMK⁺04]. After SEM inspection the titanium was then etched using $\text{HF}:\text{H}_2\text{O}$ (1:10) for 15 seconds (figure 2.11d). For the trap in chapter 8 a surface profiler⁴⁰ became available which was much more convenient for film-thickness measurements than the off-site SEM. A disadvantage is that the profiler does not give a measure of surface roughness. In this trap we required $\sim 5 \mu\text{m}$ of gold, at which thickness we were unable to repeatably produce as good a finish with some noticeable dulling of the surface. Final cleaning was carried out using acetone, IPA and then DI water. Any remaining organic residue was then removed by placing the trap under an O_3 -producing mercury lamp for several minutes.

2.5 Vacuum system

The vacuum system for traps packaged in ceramic pin grid array (CPGA) packages (as used in chapters 3-5) as built around a 114 mm diameter CF flange oc-

⁴⁰Veeco Dektak 150

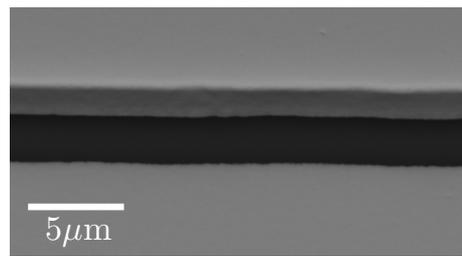


Figure 2.12: SEM (JEOL - JSM-6480) image of an inter-electrode gap taken at 55° to the plane to show the gold thickness ($\approx 2.7 \mu\text{m}$).

tagon⁴¹ containing a chip socket (see figure 2.13). The socket is bolted to a pair of goove-grabbers⁴² which lock into grooves on the inside of the octagon. The socket itself is comprised of two UHV-compatible plastic (Vespel SP-3) plates with pin receptacles⁴³ confined between them as described in [Sti07]. The receptacles for dc voltages are crimped to a 25-way ribbon cable⁴⁴ which is plugged into a D-sub vacuum feedthrough. The receptacle for the rf voltage goes via a single kapton coated cable to a 1-pin feedthrough on the octagon. Above the trap, at a distance of ~ 30 mm, is a fused-silica imaging window. It has an anti-reflection (AR) coating on the outside and a conductive indium tin oxide (ITO) coating on the inside ($4.7 \text{ M}\Omega$ from centre of glass to flange) to prevent charging. Transmission is 87% at 397 nm. The side windows are fused silica in 1.33" CF flanges with a broad-band AR coating on both sides. A list of all the materials used in the vacuum system and their preparation can be found in appendix A.

The system is pumped by an ion pump⁴⁵ and a titanium-sublimation pump⁴⁶ and pressure is measured with an ionisation gauge⁴⁷. After a 190°C bakeout the pressure reached $\sim 3 \times 10^{-10}$ Torr for the work in chapter 3. This

⁴¹Kimball Physics MCF450-SphOct-Ec2A8, 316L stainless steel with 2x 4.5" CF flanges and 8x 1.33" CF flanges

⁴²Kimball Physics MCF450-GrvGrb-C02

⁴³MILL-MAX part 0672-4-15-15-30-27-10-0, gold-plated brass shells with gold-plated beryllium copper contacts

⁴⁴Kurt J Lesker part FTAk2512P, 28 AWG kapton-coated silver-plated copper wires with monofilament PEEK weaving and a PEEK D-sub connector

⁴⁵Varian VacIon Plus 20 Diode, 271/s

⁴⁶mounted inside a nipple with internal surface area of 400 cm^2

⁴⁷Varian UHV-24P with x-ray limit of 5×10^{-12} Torr

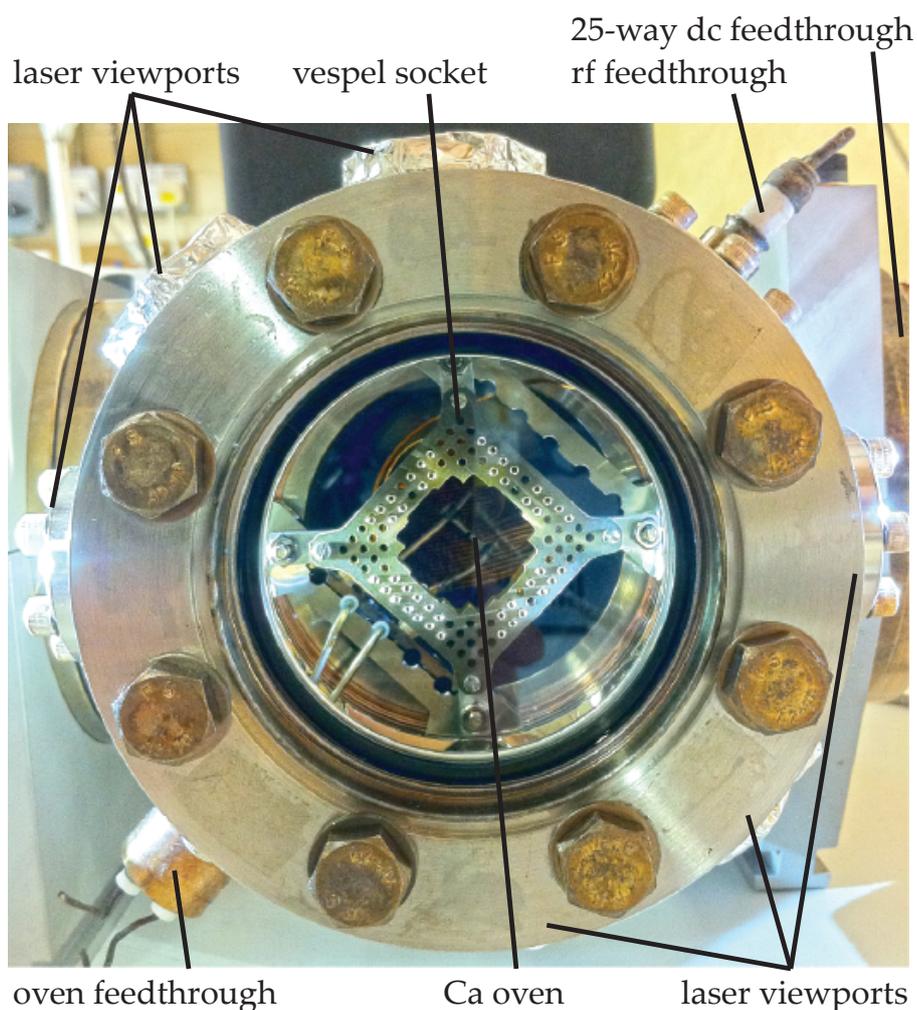


Figure 2.13: The vacuum system, as configured for the work in chapter 4. For the work in chapter 3 the oven feedthrough was located on the top flange of the octagon with the oven to the side of the socket rather than behind it.

was higher than expected and was traced to a slow leak in the all-metal valve used⁴⁸. Once this was solved the pressure reached $< 10^{-11}$ Torr for the three subsequent traps installed in the system (see chapters 4-5). Leaking valves have been an issue on several systems so we now employ a pair of valves in series so in the event of a small leak on the outer valve the inner valve can just be closed rather than changing the valve and rebaking the system.

⁴⁸Kurt J Lesker part VZCR40R

2.5.1 Calcium ovens

We produce neutral calcium beams for loading the trap by heating pieces of calcium metal in an oven. The ovens consist of a 10-15 mm length of stainless steel tube (316L stainless steel, 2.0 mm OD, 1.8 mm ID) containing the metal which is crimped at each end and has a 0.3-0.5 mm diameter hole in the side for the calcium to exit. The oven is spot-welded at the crimped sections to 2 mm diameter 316 stainless steel rod which is then connected to two pins of an electrical UHV feedthrough with barrel-connectors (gold-plated beryllium copper with stainless steel screws). Ovens were loaded immediately prior to pumping down the vacuum system (~ 15 mins being a typical time between opening calcium to air and pumping down) and doused in a stream of argon gas whenever practicable to limit oxidation.

This design of oven has been found in these experiments, and several others, to produce sufficient neutral calcium flux for rapid loading at between 4 A and 5 A. We have found though that when first used the oven has to be run briefly much hotter to obtain any flux, something we believe may be associated with 'cracking' the oxide layer on the metal pieces. The current needed to do this is between 5.5 A and 6.5 A. During this process the 423 nm resonant photoionizing laser is aligned on the trap and continually frequency scanned across the transition. Thus by monitoring the fluorescence signal on the PMT the neutral calcium flux can be monitored. This is done so that as soon as the oven starts working and the calcium flux rapidly starts increasing, the current can be reduced. In this way the trap is not exposed to a higher calcium flux than during normal operation. Degassing during this process typically causes the pressure to reach a few 10^{-8} Torr but drops enough to trap ions in a few hours and returns to its initial level after about one week. The thorough outgassing also means that at the usual operation current the oven only causes the pressure to rise by $\sim 1 \times 10^{-11}$ Torr.

Unless there is a back-side loading hole through trap (as was the case in chapters 4-5), surface traps must be loaded with the oven to the side of the trap. Alignment is achieved by ensuring there is a viewport on the opposite side of the trap to the oven. The oven is then adjusted such that the hole is just visible above the plane of the trap when looking across the surface. In this way the neutral beam is close to parallel with the trap surface to reduce calcium deposition on the electrodes and prevent shorting but it is guaranteed that the edge of the substrate is not skimming the beam and preventing loading.

2.5.1.1 Isotopic enrichment

For the work in chapter 3 natural abundance calcium metal was used. It is possible to load $^{43}\text{Ca}^+$ from a natural abundance oven by utilising the isotope shift of the resonant 423 nm transition during photoionization [LRH⁺04], however its low abundance (0.14%) limits the loading rate to only 33% of the $^{40}\text{Ca}^+$ loading rate. Loading multiple ion crystals of $^{43}\text{Ca}^+$ only can therefore be time-consuming. This is acceptable for a deep trap that can hold ions for days once loaded but is impractical for surface ion traps at room temperature that typically need reloading once an hour or so. Therefore for the work in chapter 4 onwards we switched to isotopically enriched ovens. ^{43}Ca is commercially available from Oak Ridge National Laboratory enriched to $\sim 80\%$ and typically supplied as $\sim 10\ \mu\text{m}$ thick 'target' foils. A previous oven loaded with this foil at Oxford oxidised completely on installation and failed to work. We therefore decided to use larger ($1 \times 1 \times 0.65\ \text{mm}$) pieces of calcium metal to prevent this. To reduce the cost somewhat we had Oak Ridge mix the enriched ^{43}Ca with 99.97% purity ^{40}Ca in order to produce calcium containing 12% ^{43}Ca and 88% ^{40}Ca . The pieces were then packaged into individual vials of 5-6, enough to load a single oven.

2.5.2 Magnetic field coils

Each trap has three orthogonal coils (or pairs of coils) supplied by three digitally programmable power supplies⁴⁹. These coils have around 100 turns and can provide a magnetic field of several Gauss in any direction. For work in chapter 8 we additionally have an extra set of high current coils which are described in that chapter.

⁴⁹Thurlby Thandar QL series

3

Prototype surface-electrode ion trap

A current practical challenge is the development of ion traps capable of storing and precisely manipulating a substantial number of ions [KMW02]. Surface-electrode, or planar, designs [CBB⁺05, SGA⁺04] with arbitrary electrode arrangements have been proposed as a method of fabricating large ion trap arrays. Their advantage lies in the fact that they can be easily fabricated onto the surface of a substrate in parallel using standard microfabrication techniques. This substrate region can contain control electronics and optics, allowing close integration with the ions. The work in this chapter was begun shortly after the first demonstrations of surface traps at NIST [SCR⁺06] and MIT [BCL⁺07] and was intended to develop the basic techniques of surface trap design, fabrication and operation. From this foundation we were able to conduct the more advanced experiments that form the rest of the thesis.

3.1 Trap design

3.1.1 Geometry

We selected a linear trap geometry with electrodes symmetrical about the trap's axial (\hat{z}) direction. This geometry is uncommon as the symmetry causes the radial principal axes to lie in the \hat{x} and \hat{y} directions (see figure 3.1a). If the Doppler cooling lasers are to pass across the trap's surface without striking it then they will have no projection in the \hat{y} direction, leaving that mode uncooled. Asymmetric designs [LGA⁺08, SCR⁺06] rotate these modes so they both couple to a beam parallel to the surface. We choose instead to split the central axial control electrode (between the radio frequency (rf) electrodes) to create a 'six-wire' geometry (see Figure 3.1b) [SSHM06]. We use these two electrodes and the outer segmented electrodes to place a static quadrupole (oriented at 45° w.r.t. the cardinal axes) over the cylindrically symmetric rf pseudopotential to break the degeneracy of the normal modes and ensure rotation of their axes to achieve efficient Doppler cooling (see Figure 3.2). One advantage of this approach is that it does not require one of the rf electrodes to be significantly larger than the other, allowing the use of smaller electrodes, leading to less capacitive coupling and lower losses in the rf trap drive — especially important when using a high-loss semiconductor substrate [LLC⁺09]. Symmetrical rf electrodes also simplify selection of dc control voltages in complicated arrangements, such as junctions [ABLW11]. Finally the gap between the two center electrodes directly below the ion can be used to provide optical access as has been demonstrated by the Sandia group [BEM⁺11]. This symmetric trap does have the disadvantage that it is not as tight or deep as the optimal 'four-wire' geometry [Wes08]. When this work was carried out it was not clear how the stability of the trap would be affected by the asymmetry between the dc and rf fields. This is because it will couple the equations of motion in the two orthogonal radial directions, invalidating

the standard stability analysis. A recent study by Ozakin and Shaikh [OS11] however shows that the primary stability region is robust to this asymmetry and furthermore for an angle of 45° between rf and dc quadrupoles the size of the stability region is actually increased.

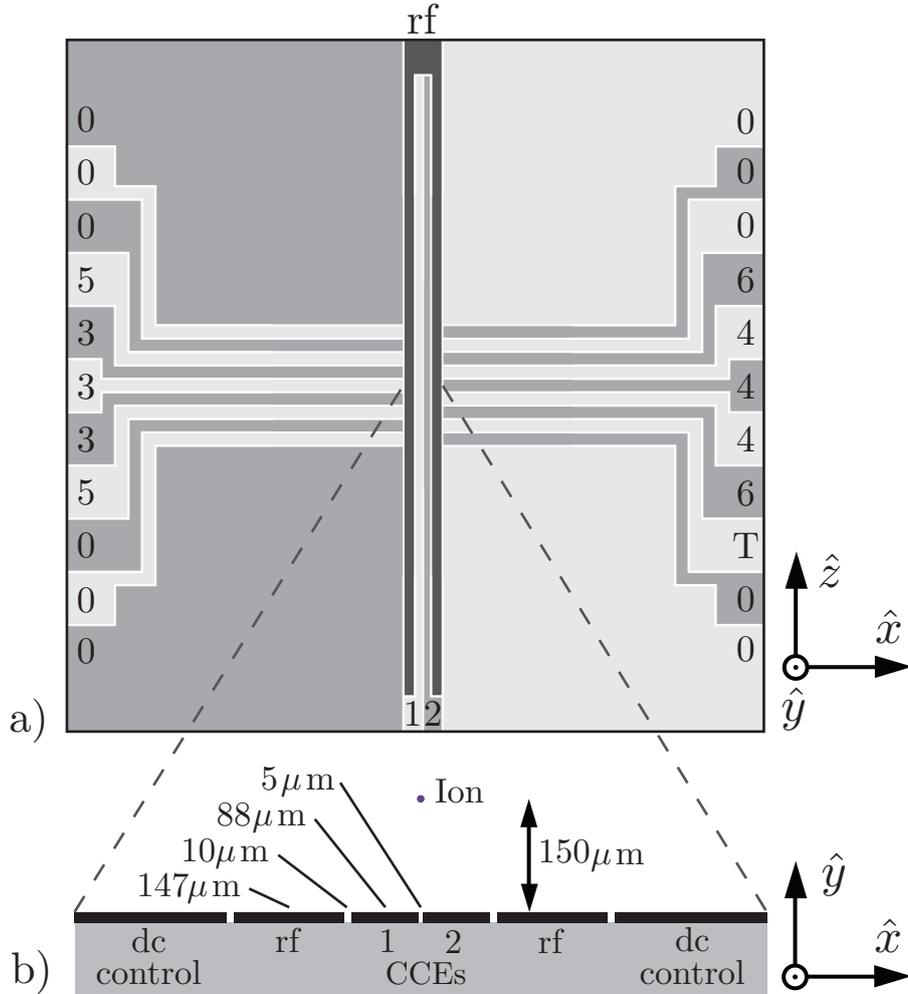


Figure 3.1: Trap design. (a) Plan view. The rf rails are dark gray and the dc control electrodes are alternately picked out in light/medium gray for clarity. The labels refer to the different voltages, $V_1 \dots V_6$, we apply. Other dc electrodes are grounded (0) except 'T' which is used for secular frequency measurements as described in Section 3.3. (b) End view showing electrode and gap widths, and ion position.

The trap has 11 pairs of outer dc control electrodes. We only control the central five pairs, grounding the others. The ions are $274 \mu\text{m}$ from the control electrodes, which are $145 \mu\text{m}$ wide in the \hat{z} direction. The ratio of electrode

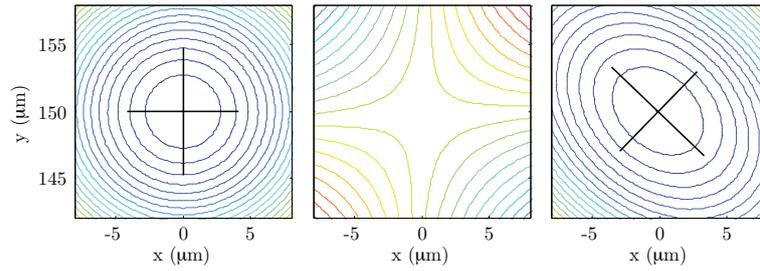


Figure 3.2: The rf pseudopotential (left), the rotation quadrupole potential (center) and the superposition of the two (right). The straight lines show the axes of the normal modes of the secular ion motion.

width to ion distance is 0.53 which is of the appropriate order of magnitude for optimal transport control [RLB⁺06].

Gaps between electrodes in the trapping region are 10 μm as these can be reliably fabricated and should give a breakdown voltage of at least 300 V [Lab08]. Directly under the ion the gap is reduced to 5 μm to reduce the exposed dielectric, as there is no high rf voltage across this gap. Far ($> 1500 \mu\text{m}$) from the ion the gaps grow to 15–25 μm to improve fabrication yield and reduce capacitive load on the rf drive. Another design feature is the extension of the rf rails across most of the chip to reduce end effects which cause axial rf field gradients, driving micromotion.

3.1.2 Operating parameters

We model our trap using numeric boundary element software (Charged Particle Optics) and verify the model using an analytical method [Wes08]. Using the model we first calculate the position and strength of the radially confining rf pseudopotential which gives us the trapped position of the ion. We then calculate the dc potential around this point due to a unit voltage applied to the i th dc control electrode. We fit a quadratic function to the calculated dc potential,

$$\Phi_i = \alpha_{xi}x^2 + \alpha_{yi}y^2 + \alpha_{zi}z^2 + \beta_{xi}x + \beta_{yi}y + \beta_{zi}z + C_i. \quad (3.1)$$

This gives us a set of quadratic α and linear β coefficients for each electrode i (the constants C_i are ignored as they do not affect the analysis). We now wish to *specify* a particular potential, for example,

$$\Phi = \alpha_x x^2 + \alpha_y y^2 + \alpha_z z^2 + \beta_x x + \beta_y y + \beta_z z \quad (3.2)$$

and find what electrode voltages $V_1 \dots V_6$ are necessary to obtain this potential.

The equation

$$\begin{pmatrix} \beta_x \\ \beta_y \\ \beta_z \\ \alpha_x \\ \alpha_y \\ \alpha_z \end{pmatrix} = \begin{pmatrix} \beta_{x1} & \beta_{x2} & \cdots & \beta_{x6} \\ \beta_{y1} & \beta_{y2} & \cdots & \beta_{y6} \\ \beta_{z1} & \beta_{z2} & \cdots & \beta_{z6} \\ \alpha_{x1} & \alpha_{x2} & \cdots & \alpha_{x6} \\ \alpha_{y1} & \alpha_{y2} & \cdots & \alpha_{y6} \\ \alpha_{z1} & \alpha_{z2} & \cdots & \alpha_{z6} \end{pmatrix} \times \begin{pmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ V_5 \\ V_6 \end{pmatrix} \quad (3.3)$$

can be used to find the required voltages. The number of independent voltages is typically fewer than six, since symmetry or Laplace's equation (which ensures $\alpha_x + \alpha_y + \alpha_z = 0$) reduces the number of degrees of freedom.

We define four sets of operating parameters, each with a required form of the potential given by the α and β coefficients in Table 3.1. The four operating parameter sets are: the 'endcap' set confining the ion axially, the 'tilt' set controlling the orientation of the radial normal modes, and sets for \hat{x} and \hat{y} micromotion compensation. The latter two apply linear fields at the site of the ion in \hat{x} and \hat{y} respectively to trim out stray fields that move the ion off the pseudopotential null and increase the driven micromotion [BMB⁺98]. To simplify operation of the trap we choose the constraints such that adjusting one parameter minimally affects the others. This is especially important in the case of the endcap set as the radial curvature it induces must be radially symmetric or it will reduce the radial normal mode tilt angle. Applying the 'tilt' potential weakly couples the \hat{x} and \hat{y} micromotion compensation. The chosen tilt voltages break the radial

mode degeneracy without overly weakening either of them or significantly reducing the trap depth. We ignore quadratic cross terms (eg. xy) in Eq. 3.1: these are zero by symmetry for the ‘endcap’ and \hat{y} -compensation bases, and negligibly small for the \hat{x} -compensation basis. For the tilt potential we use an alternate basis [$\hat{x}' = (\hat{x} + \hat{y})/\sqrt{2}$, $\hat{y}' = (\hat{x} - \hat{y})/\sqrt{2}$ and $\hat{z}' = \hat{z}$] so that again the quadratic cross-terms vanish by symmetry. In this model the tilt angle would always be 45° as soon as the radial mode degeneracy was lifted; in practice, end effects modify this angle. The angle can be estimated (see Table. 3.3) from a fit to the full potential (sum of rf pseudopotential, ‘endcap’ potential and ‘tilt’ potential).

The endcap voltages shown in Table 3.2 give an axial secular frequency of 500 kHz. The radial secular frequency is independently set by the rf voltage applied to the rf electrodes. At our typical rf voltage, the ‘tilt’ voltages in Table 3.2 give radial secular frequencies of 2.29, 3.33 MHz and an estimated tilt angle of 42° .

	β_x	β_y	β_z	α_x	α_y	α_z
Endcap	0	0	0	-1.02×10^6	-1.02×10^6	2.05×10^6
\hat{x} -Comp	1	0	0	0	0	0
\hat{y} -Comp	0	1	0	0	0	0
	$\beta_{x'}$	$\beta_{y'}$	$\beta_{z'}$	$\alpha_{x'}$	$\alpha_{y'}$	$\alpha_{z'}$
Tilt	0	0	0	1.00×10^7	-1.00×10^7	0

Table 3.1: The coefficients of the potentials for our four voltage bases. β is in units of Vm^{-1} and α in units of Vm^{-2} . For the ‘tilt’ parameters we use the alternate basis $\hat{x}' = (\hat{x} + \hat{y})/\sqrt{2}$, $\hat{y}' = (\hat{x} - \hat{y})/\sqrt{2}$ and $\hat{z}' = \hat{z}$.

	V_1	V_2	V_3	V_4	V_5	V_6
Endcap	-1040	-1040	-3152	-3152	-235	-235
\hat{x} -Comp	0	0	-0.95	0.95	-0.95	0.95
\hat{y} -Comp	0.92	0.92	1.86	1.86	5.02	5.02
Tilt	-886	929	1117	-1030	1117	-1030

Table 3.2: The calculated voltage bases (in mV).

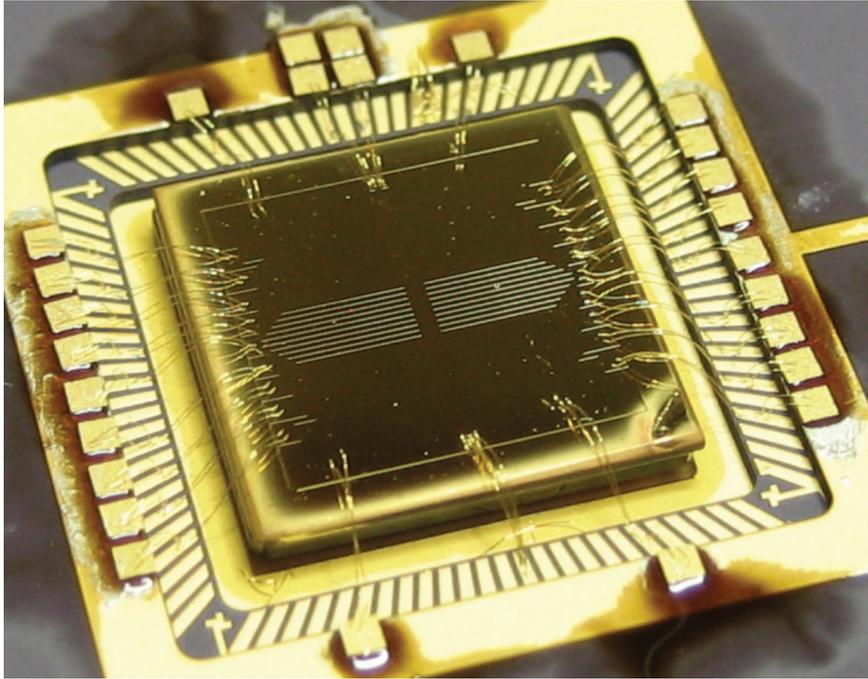


Figure 3.3: Final packaged trap. The gold squares are the single layer capacitors.

3.2 Experimental apparatus

3.2.1 Trap packaging

The trap is fabricated on a $10 \times 10 \times 0.5$ mm polished single-crystal quartz substrate. This was glued to a 100-pin ceramic pin grid array (CPGA) package¹ with Epo-Tek 353ND epoxy via a 1mm fused silica spacer (see figure 3.3). Each electrode is wire bonded to the package with two bonds except the rf which has five bonds. 820 pF single layer capacitors² are mounted directly to the package to provide a low-impedance path to ground for any rf coupled onto the dc electrodes. These are attached to the outer gold ring of the CPGA using Epo-Tek H20E conducting epoxy. 353ND epoxy is then applied around each capacitor to prevent them being shorted by calcium deposition from the oven used to load the trap. The outer ring is grounded and the capacitors wire-bonded to each

¹Kyocera drawing number IPKX0F1-8180AB

²American Technical Ceramics part number 116UL821M100TT

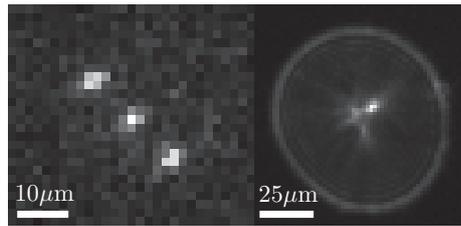


Figure 3.4: Three ions in the trap (left) and a single ion with the imaging system focused on the *image* of the ion in the electrodes (right).

electrode's pad on the package with two bonds. Two capacitors are used on each of the centre control electrodes as they have higher capacitance to the rf rails.

3.2.2 Voltage supplies

The rf trap drive is produced by a Stanford Research Systems DS345 Synthesizer, amplified by a Mini-Circuits ZHL-1-2W, passed through a directional coupler (Mini-circuits ZDC-20-3) for monitoring VSWR and finally stepped up with an inductively coupled helical resonator. The loaded resonant frequency of the resonator is 25.8 MHz. A two-turn pickup coil around the rf vacuum feedthrough allows us to monitor the trap voltage. Dc voltages are calculated using a Lab-view program which takes desired tilt, axial frequency and micromotion compensation parameters and produces the required voltages with a 16-bit DAC (Measurement Computing PCI DAC6703). The output from the DAC is filtered using a C-R-C II filter ($R=1\text{ M}\Omega$, $C=0.1\text{ }\mu\text{F}$) just before the dc vacuum feedthrough.

3.3 Trapping results

Up to three ions (see figure 3.4) have been successfully trapped and crystallized. A ratio of fluorescence signal to background scatter of over 100 is observed for a single ion. We can focus the imaging system on either the ion itself or on its

reflection in the gold electrodes. The reflection increases the fluorescence collection efficiency by up to $\approx 50\%$ compared to a non-surface-electrode trap studied previously. This reflection reduces the readout error for a single ion [MSW⁺08] but also significantly increases cross-talk from neighbouring ions when reading out one ion from a multiple ion string. Based on simulations using a maximum likelihood method [Bur10] we estimate the readout error caused by this crosstalk can be reduced to the level of 10^{-3} . If higher readout fidelity is required then a spatial filtering scheme or imaging from a different direction could be employed.

3.3.1 Ion lifetime

Ion lifetime is around 5 minutes with cooling lasers on and around 15 seconds with lasers off. It is thought to be limited by collisions resulting from the relatively high background pressure (3×10^{-10} Torr). Similar issues with low lifetime in shallow surface traps at $\approx 10^{-10}$ Torr pressures have been reported elsewhere [LLC⁺09]. Numerical simulations of two-dimensional (in the radial plane) ion trajectories were carried out to investigate the possibility of collisional loss. Figure 3.5 shows the calculated probability of immediate loss (within $2 \mu\text{s}$ of collision) of the ion versus the initial kinetic energy imparted to it by a collision, assuming the ion was at rest before the collision. Each point is averaged over different rf phases and collision angles. The results show that ion loss is possible at energies of about half the trap depth given by the calculated pseudopotential. However, the rate of ‘hard’ collisions (where background neutrals penetrate the angular momentum barrier to collision), as calculated using Langevin theory [WMI⁺98a], which impart the required collision energy is not enough to explain the observed loss rate. One possible loss mechanism is that collisions, including more frequent ‘soft’ collisions (which do not penetrate the angular momentum barrier), put the ion into a highly excited orbit where large Doppler shifts and poor laser beam overlap reduce cooling efficiency. The heating rate for a hot

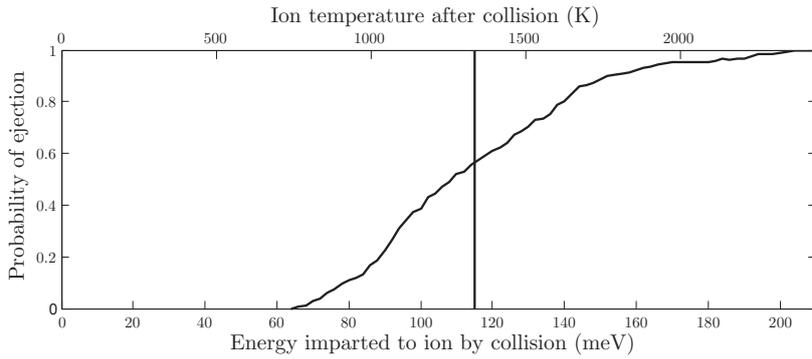


Figure 3.5: Computed probability of immediate ion loss from a collision. The vertical line is the depth of the pseudopotential.

($\gtrsim 100$ K) ion will also be greater than we measure in Section 3.5 (where we measure ion temperatures of up to ≈ 1 K), as away from the rf null the ion is heated by noise on the rf supply and the large anharmonicities in a surface trap away from its center will cause micromotion to heat the ion [WMI⁺98a]. These effects may cause the ion to be heated out of the trap more quickly than it is cooled. This is corroborated by the fact that a doubling of trap depth does not give a significant increase in trapping lifetime. In this model the increased trap depth could be largely cancelled by the increased micromotion (see Eq. 3.4) and increased heating rate (see figure 3.12). We also see drop-outs in fluorescence of several hundred milliseconds every minute or so which could be due to collisions (they occur with the 854 nm repumper on so are not quantum jumps to the $D_{5/2}$ shelf).

3.3.2 Secular frequencies

The experimental and modelled radial secular frequencies as a function of the applied radial mode tilt are shown in Table 3.3. These are found experimentally by applying an oscillating voltage to an electrode (marked ‘T’ in figure 3.1) which observably heats the ion when resonant with a motional mode. The axial secular frequency was 467 kHz compared to the model value of 500 kHz. By working the secular frequencies and the rf drive frequency (25.8 MHz) back

Tilt factor	$f_{expt.}$ (MHz)	f_{model} (MHz)	θ_{model}
0	3.11, 3.16	3.12, 3.15	0°
0.125	3.11, 3.18	3.10, 3.16	28°
0.25	3.09, 3.19	3.08, 3.19	35°
0.5	3.03, 3.24	3.03, 3.24	40°
1	2.90, 3.34	2.92, 3.33	42°
2	2.63, 3.52	2.70, 3.52	44°

Table 3.3: Comparison of measured and modelled radial trap frequencies for different tilt factors. θ_{model} is the angle of tilt for the normal mode initially perpendicular ($\theta = 0^\circ$) to the trap plane. A tilt factor of 1 corresponds to $\alpha_{x'} = 1.0 \times 10^7 \text{ Vm}^{-2}$ (see Table 3.1.)

through our model of the trap we calculate an rf amplitude of 175 V, a stability parameter of $q=0.34$ and a trap depth of 115 meV. We managed to operate the trap with rf amplitude between 223 V ($q=0.43$, $f_{radial} = 4.02$ MHz, depth=188 meV) and 112 V ($q=0.22$, $f_{radial} = 1.99$ MHz, depth=47 meV). The rf amplitude is the only model parameter that is fitted to the data. The voltage amplitude before the resonator was 7.07 V implying a resonator step-up of 24.7.

3.4 Micromotion compensation

In a linear Paul trap any displacement of the ion off the rf null in the radial directions (\hat{x} or \hat{y}) of the trap will lead to the ion experiencing driven motion, known as micromotion, due to the confining rf field [BMB⁺98]. The expressions for the radial displacement x_d of the ion from the rf null due to field E_{dc} and for the amplitude x_μ of the resulting excess micromotion are

$$x_d = \frac{QE_{dc}}{m\omega_r^2}, \quad x_\mu = \sqrt{2} \frac{\omega_r}{\Omega_{rf}} x_d, \quad (3.4)$$

where ω_r is the radial trapping frequency, Ω_{rf} is the trapping rf frequency, m is the ion's mass and Q is the ion's charge. This motion is undesirable as it will cause Doppler shifts, allows noise on the rf supply to couple to the ion and in certain circumstances can directly heat ions [WMI⁺98a]. A widely used method of detecting micromotion is to measure the correlation of the arrival time of photons scattered during Doppler cooling with the rf phase [BMB⁺98].

This correlation arises due to first order Doppler shift from the micromotion changing the apparent detuning of the cooling lasers at different points in the rf cycle. It is observed by recording a histogram of the output from a time-to-amplitude converter which measures the delay between photon arrival and a sync pulse from the trap rf supply (see section 2.2.1.1). The excess micromotion can then be reduced by suitably modifying the static voltages applied to the trap's electrodes to shift the ion back to the rf null.

3.4.1 Out-of-plane micromotion detection

Doppler cooling beams must lie in the $(\hat{x} - \hat{z})$ plane of surface electrode traps to avoid striking the trap and causing scatter or charging of the substrate. Full radial micromotion compensation cannot be carried out in this geometry using the standard correlation technique. This is because there will necessarily be a radial direction in which the ion can move (\hat{y}) which is perpendicular to any laser direction and will therefore not have a Doppler shift associated with it.

Ions with a metastable low-lying D state require a repumping laser to clear out population that decays into it during Doppler cooling. Here we demonstrate that precise rf correlation in all radial directions can be carried out with this repumping laser placed in an out of plane direction (see figure 3.6). This will cause the laser to strike the surface of the trap (and be mostly reflected by the metallic electrodes). Since this repump laser is in the infrared, unlike the ultraviolet Doppler cooling laser, it is unlikely to cause photoemission and charging of the trap substrate (we confirmed this experimentally; see section 3.4.4 below). Due to the low branching ratio (6%) of the repumper transition we cannot rely on differing repumping rates at different Doppler shifts to give us a correlation signal. Instead we operate with high repumper intensity ($I_r > 50I_s$ where $I_s \equiv 4\pi\hbar c\Gamma/\lambda^3$) where a coherent Raman process allows the repumper's Doppler shift to modulate directly the $P_{1/2}$ population, and thus the fluores-

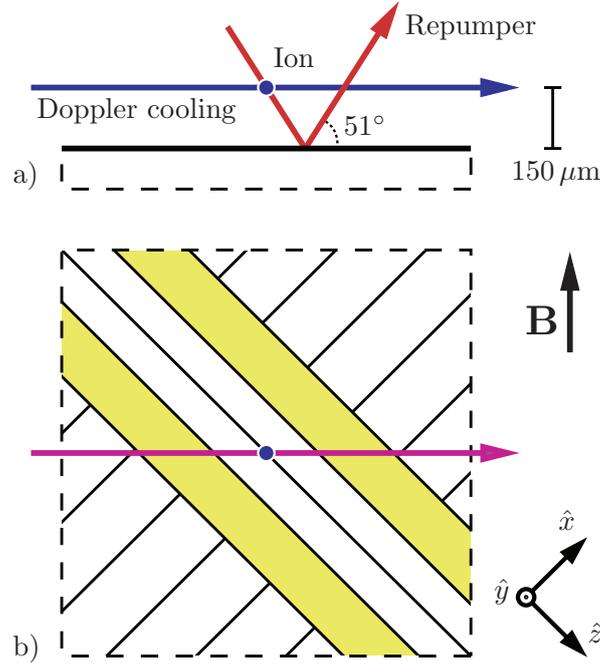


Figure 3.6: Beam directions used during \hat{y} micromotion compensation. For \hat{x} micromotion compensation the repumper is co-linear with the Doppler cooling beam.

cence.

3.4.2 Theoretical modelling of modulated Raman effect

To obtain the optimal parameters for this technique we analyzed the Bloch equations with the inclusion of laser frequency modulation following the method in [Obe99]. The equation for the density matrix of atomic states $\tilde{\rho}$ is given by

$$\frac{d\tilde{\rho}}{dt} = (M_0 + \Delta M \cos(\Omega t))\tilde{\rho} \quad (3.5)$$

where the matrix M_0 is the 8×8 Liouvillian describing the Bloch equations of our $S_{1/2}$ - $P_{1/2}$ - $D_{3/2}$ system (including Zeeman structure) and Ω is the trap rf drive frequency. $\Delta M \equiv 2\pi v_0 \Delta M_r / \lambda_r$ where λ_r is the repump laser's wavelength, v_0 is the velocity amplitude in the direction of the laser beam and ΔM_r is the linear change in M_0 with unit repumper detuning. The steady state solution of this

equation can only contain frequency components at multiples of Ω and so takes the form

$$\tilde{\rho} = \sum_{n=-\infty}^{\infty} \tilde{\rho}_n e^{-in\Omega t}. \quad (3.6)$$

Substituting Eq. 3.6 into Eq. 3.5 gives the recursion relation

$$(M_0 + in\Omega)\tilde{\rho}_n + \frac{1}{2}\Delta M(\tilde{\rho}_{n+1} + \tilde{\rho}_{n-1}) = 0 \quad (3.7)$$

Terms of $|n| > 1$ are negligible in the limit of $\frac{k_r v_0}{\Omega} \ll 1$ and we obtain solutions determined by

$$(M_0 - \frac{1}{2}\Delta M(M_0^2 + \Omega^2)^{-1}M_0\Delta M)\tilde{\rho}_0 = 0 \quad (3.8)$$

$$\tilde{\rho}_{\pm 1} = -\frac{1}{2}(M_0 \pm i\Omega)^{-1}\Delta M\tilde{\rho}_0 \quad (3.9)$$

which can be evaluated numerically and substituted back into Eq. 3.6 to give the mean fluorescence and its modulation due to micromotion. We assume laser linewidths of 500 kHz, σ^\pm polarizations, a B-field of 1.7 G and ignored $D_{3/2} \rightarrow S_{1/2}$ decay. A systematic search of the parameter space of laser intensities and detunings for both the Doppler cooling and repump lasers was carried out to give the optimal parameters. Figure 3.7 shows the maximum micromotion sensitivity as a function of Doppler cooling laser detuning Δ_c and repump laser intensity I_r . Increasing the Doppler cooling intensity I_c causes the peak sensitivity to drop off $\propto (1 - 0.07I_c/I_s)$ but we fix it at $I_c = 1.5I_s$ as this was found experimentally to be a good trade off between sensitivity and the increased signal and cooling given by a higher I_c . The repumper detuning Δ_r is optimized at each point and is found to be in the range of 10–20 MHz to the red of the Doppler cooling laser. This is fortuitous as it avoids the ion heating that can occur when $\Delta_r > \Delta_c$ even for $\Delta_c < 0$ (see figure 3.8).

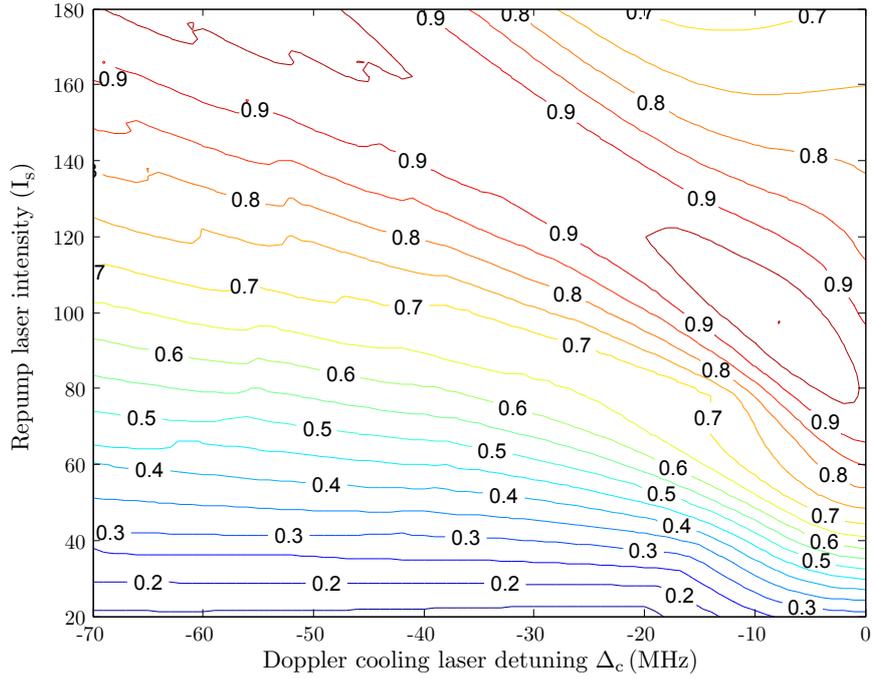


Figure 3.7: Relative sensitivity of micromotion compensation for $I_c=1.5I_s$ and Δ_r optimized at each point.

3.4.3 Experimental implementation

In the limit of large micromotion the ion sees most of the laser power in motional sidebands which gives a correlation signal that is hard to interpret as the model above is no longer applicable. Therefore we first coarsely compensate the out-of-plane micromotion by looking at profiles of fluorescence versus Δ_r , a technique described in [LKC⁺05]. We find that with $E_{dc} \lesssim 100 \text{ Vm}^{-1}$ the scan gives a well resolved single dark resonance and we can proceed with the correlation method.

Using this technique the maximum \hat{y} -direction sensitivity was a relative fluorescence modulation of 0.9(1)% per Vm^{-1} (see figure 3.9) with the laser parameters as in figure 3.8. In comparison we obtain \hat{x} -direction sensitivity of 1.3(1)% per Vm^{-1} using the typical co-linear beam arrangement and the Doppler laser detuned to the half-fluorescence point for maximum sensitivity. In practice we typically trim stray fields to below 1 Vm^{-1} in \hat{x} and 3 Vm^{-1} in \hat{y} . This corresponds to peak velocity components of 0.1 m/s and 0.3 m/s respectively.

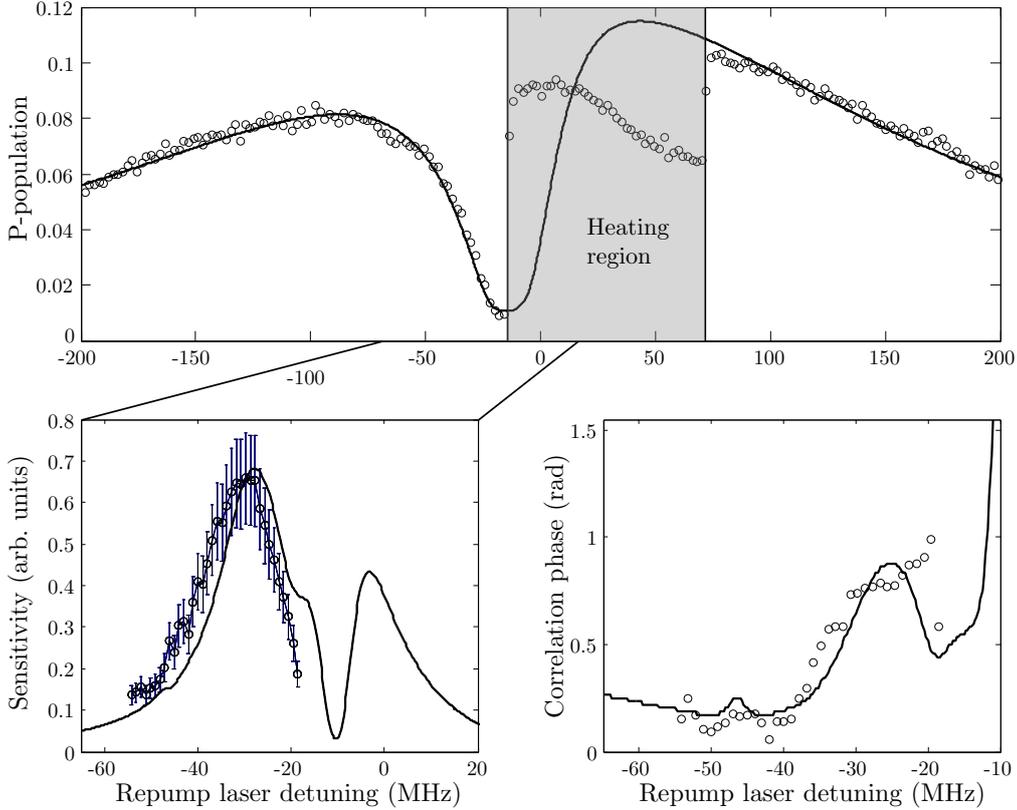


Figure 3.8: Experimental 866 nm repumper scan (top) and model fitted to give $I_c = 1.7I_s$, $I_r = 95I_s$ and $\Delta_c = -14$ MHz. Predicted micromotion sensitivity based on laser parameters and the experimentally observed sensitivity (bottom left) with vertical scale fitted to the data. Predicted correlation phase and the experimentally observed phase with the unknown offset fit to the model (bottom right).

3.4.4 Micromotion compensation results

Without the tilt potential the rf rail symmetry would mean an \hat{x} (\hat{y}) compensation field would null \hat{x} (\hat{y}) micromotion. The tilt voltages introduce some cross-coupling (of order 10%) so iterations between \hat{x} (performed in the standard manner) and \hat{y} compensation are needed.

We detect no significant uncompensatable micromotion which would suggest rf phase mismatch [BMB⁺98] or axial micromotion due to an rf-field gradient along \hat{z} .

We observe micromotion compensation fields of up to 1000 Vm^{-1} in both directions with drifts on the order of 10 Vm^{-1} per hour and no noticeable step-

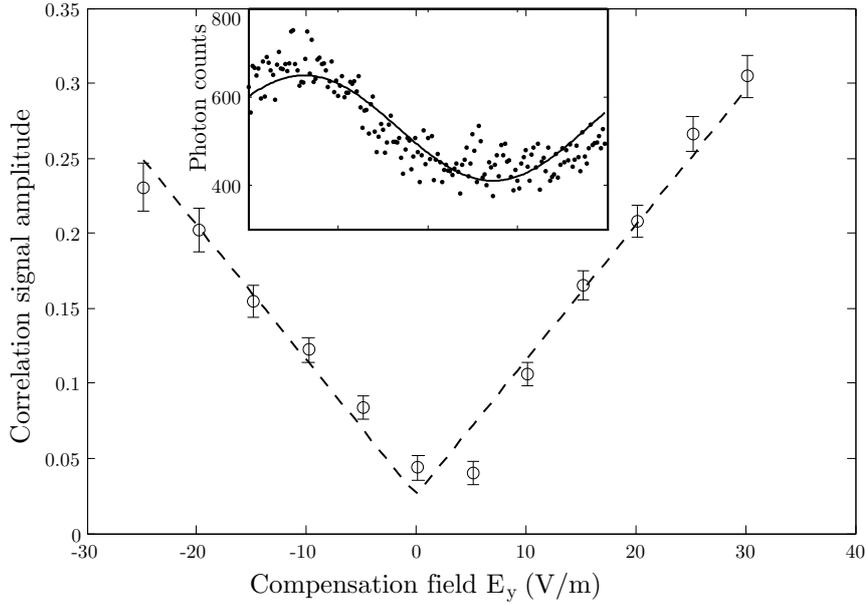


Figure 3.9: Fitted relative amplitude of rf correlation for different \hat{y} compensation settings and a sample set of correlation data (inset). Laser parameters are as fitted in figure 3.8 with $\Delta_r = -28.7$ MHz

change during loading. In light of this drift rate, a more accurate compensation technique is not required. This fairly low drift rate is possibly due to the fact the ion only has direct line of sight to two regions of exposed dielectric, that of the insulating gaps between electrodes and that of the small laser windows. The former is shielded somewhat by the aspect ratio of the gold electrodes and the latter by the vacuum apparatus and the trap structure. To check that the out-of-plane 866 nm repump beam does not cause a change in the micromotion we increased its power to 1 mW ($\approx 300I_s$) and translated it approximately 100 μm off the ion in the direction of the in-plane beams. We then monitored the \hat{x} compensation using those beams whilst alternately turning the out-of-plane beam on and off for several minutes at a time. There was no detected correlation. For Doppler cooling we typically use a 397 nm intensity $\simeq I_s$ (corresponding to $\simeq 5 \mu\text{W}$), and 20 μW of 389 nm, 80 μW of 423 nm light were used for photoionization loading.

3.4.5 Comparison with other schemes

Other methods of compensating micromotion in the out-of-plane direction exist. One is to measure the ion position as a function of rf voltage by translating the Doppler cooling laser [BCL⁺07]; however this is time-consuming and not particularly sensitive. Another method is to add a voltage component to the rf drive at $\Omega_{rf} \pm \omega_r$ to force parametric heating of the ion, typically by scanning the frequency of the added voltage and observing a dip in fluorescence on resonance. As this parametric heating effect does not occur when the ion is on the rf null, compensation can be achieved by minimising this dip [ITU11, SDM⁺11]. This method has two main disadvantages compared to ours. Firstly it does not give a quantitative measurement of the residual micromotion velocity present, just a 'best compensation field'. Secondly, our method has an accuracy limited only by integration time whereas this method will require ever larger noise fields to increase sensitivity which is impractical as any brief perturbation, such as a collision, will move the ion off the null and cause rapid heating. It does have the important advantage over our method in that it is applicable to ions without low-lying D-states which are in use by many groups. It also does not require a laser to strike the trap, which was not a problem here, but may be undesirable in some experiments or in the case of Ba⁺ where the repumper has a much shorter wavelength (650 nm).

3.5 Heating rate measurement by Doppler recooling

3.5.1 Modified scheme for Lambda systems

For projective measurement of the S_{1/2}-D_{5/2} optical qubit [MSW⁺08] or micromotion compensation we repump decays into the D_{3/2} level using the 866nm transition (see figure 3.10a). Coherent dark resonance effects and stimulated emission into the D_{3/2} state complicate the analysis of some experiments such

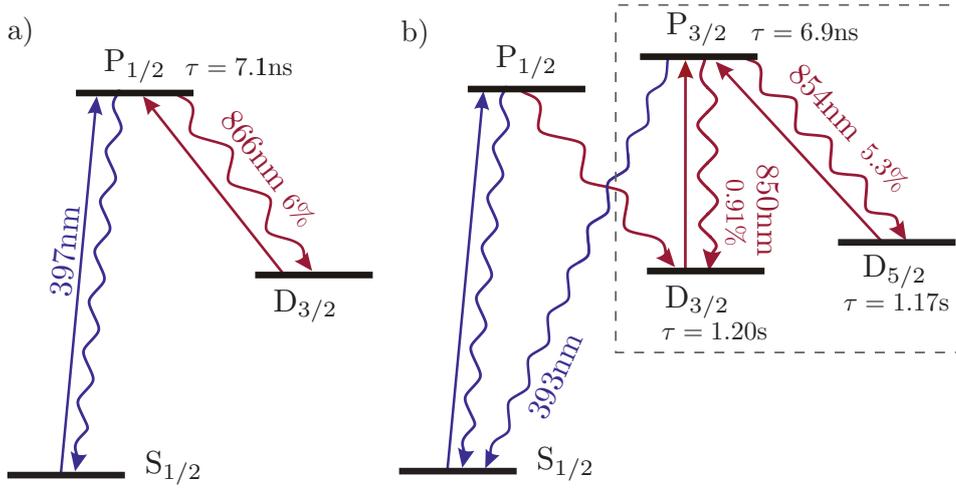


Figure 3.10: Doppler cooling with (a) 866 nm repumping and (b) 850 nm/854 nm repumping. The percentages are the branching ratios of the three repump transitions. In (b), because there is no laser linking the boxed level system with the $S_{1/2}$ - $P_{3/2}$ system, and the $P_{1/2} \rightarrow D_{3/2}$ branching ratio is small, the $S_{1/2}$ - $P_{1/2}$ system behaves like a quasi two-level atom.

as heating rate measurements and limit the fluorescing $P_{1/2}$ population to $1/4$. However, if the $P_{3/2}$ - $D_{5/2,3/2}$ transitions at 850 nm and 854 nm are used to pump the 6% decay to $D_{3/2}$ back to the $S_{1/2}$ state (see figure 2.2b) these effects are avoided. (There is still a dark resonance if the 854 nm and 850 nm detunings match but for co-propagating beams the difference in the Doppler shifts is small enough that if initially set to different detunings they will not be shifted into a dark resonance by ion motion.) We apply high intensity ($I \sim 10^3 I_s$) 850 nm and 854 nm light to power-broaden these transitions such that the repump rate is broadly insensitive to Doppler shift; the $S_{1/2}$ - $P_{1/2}$ Doppler cooling transition acquires a largely Lorentzian lineshape with deviations due to a small population in the $D_{3/2}$ - $P_{3/2}$ - $D_{5/2}$ manifold. Figure 3.11 shows fluorescence data fitted by a two-level model where the excited state population is given by

$$\rho_{ee} = \frac{s/2}{1 + s + (2\Delta/\Gamma)^2} \quad (3.10)$$

where s is the saturation parameter $s \equiv 2|\Omega_{Rabi}|^2/\Gamma^2$, Δ is the Doppler cooling laser detuning and Γ is the effective linewidth, broadened by laser linewidth and the Zeeman splitting. The close agreement of the data to a Lorentzian lineshape

means we can interpret the Doppler recool method heating rate measurements, as described in [WEL⁺07] for a two-level system.

Since this work was carried out, the Berkeley group have developed an alternative method [DNM⁺11] for adapting the Doppler recool method to Ca⁺ (and other ions with low-lying D-states). Their scheme is to use 397 nm and 866 nm cooling and then analyse it using the standard scheme in [WEL⁺07]. They then rescale the incorrect heating rate thus obtained using a calibration obtained by injecting white noise onto one of the trap electrodes. This method has several disadvantages. First the recalibration has to be done every time there is a chance the laser parameters may have drifted (they report a factor ~ 2 change in calibration under nominally identical conditions). Secondly it is hard to predict what the actual amplitude of an externally injected noise source will be at the ion due to stray inductance and capacitance throughout the system, particularly when dc filters are included in the vacuum system. It does have the advantage that it does not require the 850 nm and 854 nm lasers, but in our case they are also required for readout (see section 6.5.0.2).

After an ion heating period of one second we measure a mean axial vibrational quanta of $\bar{n} \approx 5 \times 10^4$ which corresponds to an electric field spectral noise density of $S_E \approx 2 \times 10^{-10} \text{ V}^2\text{m}^{-2}\text{Hz}^{-1}$ (at 3.1 MHz) which is comparable to traps of similar size in the literature [SCR⁺06]. Before this data was fully acquired, a fault in the rf supply caused the trap to arc. This arcing visibly damaged the surface quality and caused electron emission points to appear at our trap operating voltages. This emission can be quantified by measuring the difference in background counts with and without the rf supply on, using a CCD camera. It shows a rapid increase above $\sim 150 \text{ V}$ and corresponds with a clear increase in the heating rate (see figure 3.12). The effect of the decreased surface quality is unknown. The heating rate also shows significant day to day shift.

3.5. Heating rate measurement by Doppler recooling

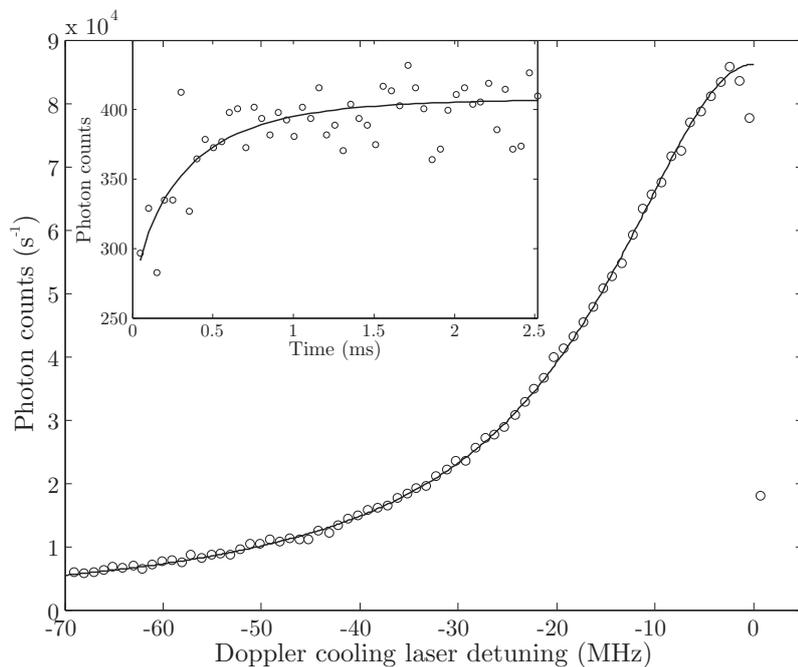


Figure 3.11: Fitted fluorescence profile of Doppler cooling laser ($s=1.04$, $\Gamma=25.5$ MHz) and (inset) a typical Doppler recool experiment ($s=1.03$, $\Delta=17.0$ MHz).

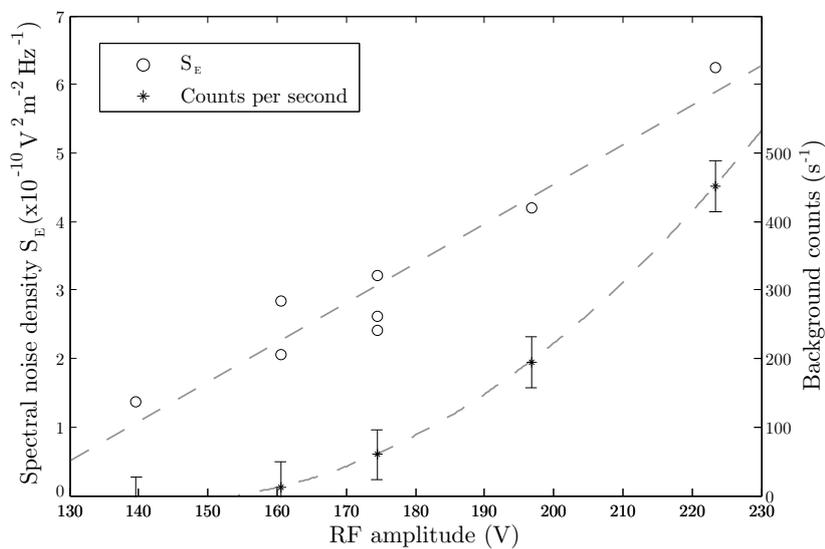


Figure 3.12: Heating rate (spectral noise density of electric field, S_E) and background camera counts due to electron emission 'glow' from electrodes, both as a function of rf trap voltage. The linear and quadratic fits shown are to guide the eye only.

4

Characterisation of a microfabricated trap

4.1 The traps

In this chapter we investigate the performance of a trap which is constructed using intrinsically scalable semiconductor fabrication technology. The fabrication process is described in more detail in [SFH⁺10], where initial characterisation of the trap is also reported. This design of trap is compatible with integrated microfabricated optical elements, and it has already been used to demonstrate collection of ion fluorescence using diffractive micro-optics [BEM⁺11]. Ion shuttling through junctions has also been demonstrated in traps utilising the same fabrication process [MHS⁺11].

We tested three different traps of the same aluminium-on-SiO₂ construction, all with an ion-to-electrode distance of 98 μm (see fig. 4.2). The traps were nominally identical to those described in [SFH⁺10] except for Trap 2 which had 13 μm high oxide pillars supporting the electrodes rather than 20 μm high pillars (see fig. 4.1). Traps 1, 2 and 3 differed in the filtering of the dc control electrodes (see section 4.5).

The traps were mounted in the same vacuum system described in section 2.5. The system was modified such that the neutral calcium oven was behind, rather

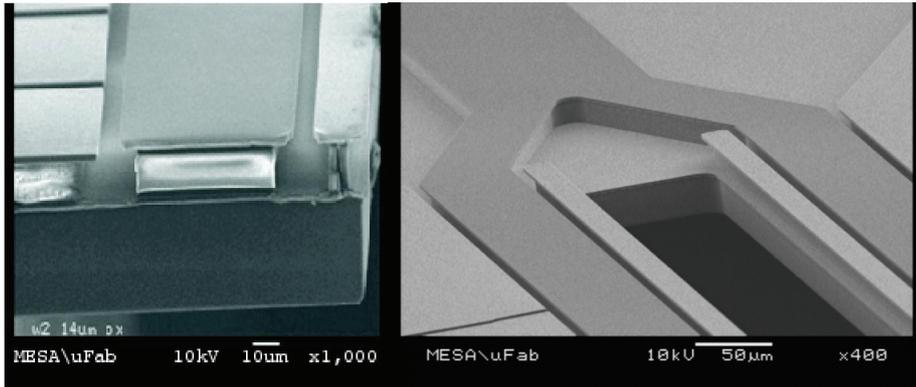


Figure 4.1: SEM image (left) showing the 5 micron overhang from the supporting oxide pillar and the 7 micron gap between neighbouring electrodes. Figure taken from [SFH⁺10]. SEM image (right) showing detail of the slot and the rf electrode where it splits into two on either side of the ion loading slot.

than to the side of, the trap so the ions are loaded through the central slot (see figure 2.13). No special precautions were taken to prevent a native oxide layer forming on the aluminium electrodes between fabrication and installation.

4.2 Voltage supplies

4.2.1 RF supply

The rf source we used was an Agilent 8247C synthesiser. A 20 dB coupler after the amplifier picks off a fraction of the power for the photon correlation electronics (see section 2.2.1.1). The rf is then amplified by a 1.58 W amplifier¹. The amplifier is followed by a bandpass filter (see below) and then finally the trapping voltage was stepped-up using a toroid. This section after the amplifier is constructed from short rigid BNC I-pieces and elbows for stability. The toroid is iron powder (Micrometals T94-6) and is in a resonant transformer arrangement (see fig. 4.5). The toroid gave a loaded Q of 28 and a resonant frequency of $\Omega_{\text{rf}} = 2\pi \times 33$ MHz. The actual voltage step-up at the trap was 21.8, which was calculated by measuring an ion's radial secular frequencies ω_r and deducing the

¹Motorola CA2842C, 22.5 dB gain, noise figure < 4 dB

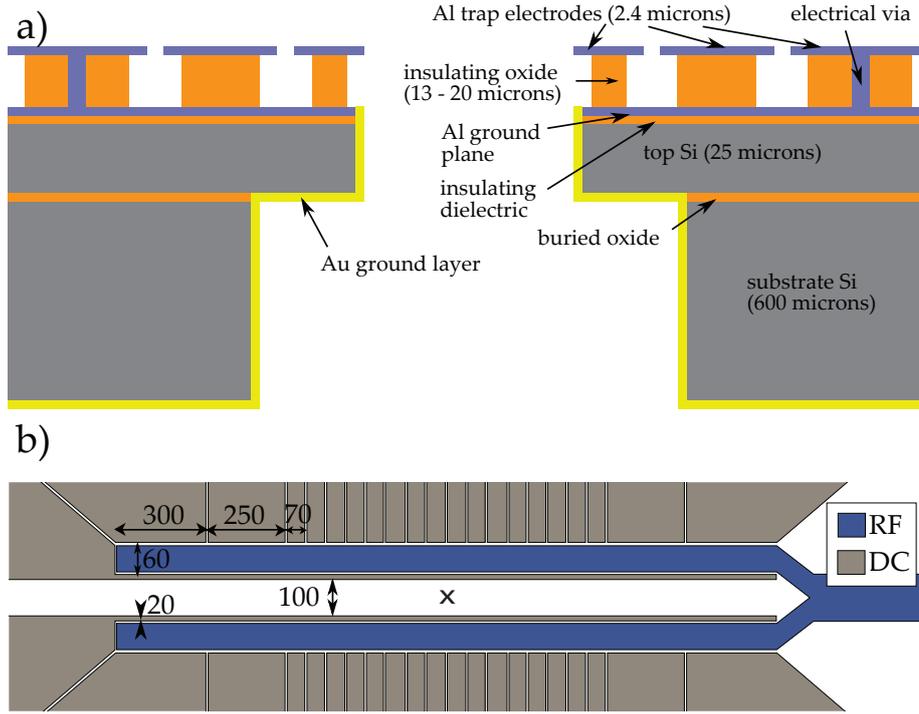


Figure 4.2: Schematic diagram of the ion trap. (a) Cross-sectional view (not to scale), with layer thicknesses indicated in microns. The ion is trapped $84 \mu\text{m}$ above the plane of the trap electrodes. (b) Plan view (to scale), with dimensions shown in microns. A \times marks the position of the trap centre used in these experiments.

trap voltage using our electric field simulation. The voltage amplitude used was in the range 50–140 V (zero-peak). A toroid was used, rather than the helical or coaxial quarter-wave resonators typically used with many ion traps [SSWH11], as they are experimentally much more convenient work with. Toroids are lossier so the voltage step-ups are typically less but for small traps requiring ~ 100 V this is not a problem.

4.2.1.1 RF filter

As the toroid does not have a high enough Q to filter out noise effectively at $\Omega_{\text{rf}} \pm \omega_{\text{r}}$ (which can heat the ion [WMI⁺98b]) we placed a bandpass filter (see fig. 4.3) between the amplifier and the toroid. The filter is a 50Ω impedance mesh-capacitor topology² with 3.2 dB insertion loss at 33 MHz, -38 dBc at

²Designed using Elsie design program from Tonne Software

4. CHARACTERISATION OF A MICROFABRICATED TRAP

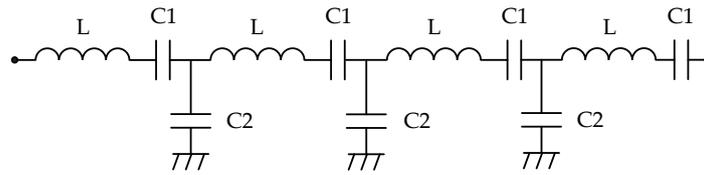


Figure 4.3: Circuit diagram and photo of the bandpass filter. The inductors are Delevan 4470 series iron core with $L=22\ \mu\text{H}$, $Q=130$ (at 10 MHz) and an SRF of 92 MHz. C1 is a Johanson 5201 air-spaced trimmer (0.8-10 pF) in parallel with a 4.7 pF ceramic leaded capacitor. C2 is a 47 pF and a 100 pF ceramic leaded capacitor in parallel.

30 MHz and $-30\ \text{dBc}$ at 36 MHz (typical radial frequencies are 3–4 MHz). The key to obtaining a low insertion loss is to use the highest Q inductors available and to use stable trimmer capacitors so that each section in turn can be tuned exactly to resonance.

4.2.2 DC supply

DC electrode voltages of up to $\pm 10\ \text{V}$ were provided by a single multichannel DAC chip³ programmed over USB. A USB isolator⁴ was placed immediately before the DAC to block noise from the PC. The voltage reference for the DAC was supplied by a low noise reference chip⁵. The DAC chip has an intrinsic voltage

³Analog Devices AD5372 32-ch, 16-bit serial input DAC on EVAL-AD5370/2/3EBZ evaluation board with USB to serial interface

⁴Analog Devices ADUM4160 on EVAL-ADUM4160 evaluation board

⁵Linear Technology LTC6655BHMS8-5 with a $10\ \mu\text{F}$ film capacitor on the output, $2\ \text{ppm}/^\circ\text{C}$ drift, $\pm 0.025\%$ accuracy, $< 20\ \text{nV}/\sqrt{\text{Hz}}$ noise above above 100 kHz

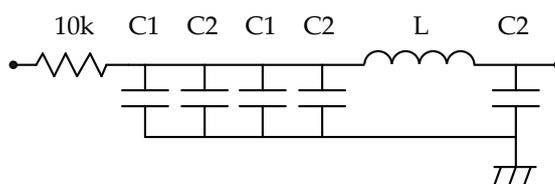


Figure 4.4: Circuit diagram of the dc filters. The inductor is a Coilcraft 1812CS ceramic chip with $L=22\ \mu\text{H}$ and an SRF of 40 MHz. The capacitors are Murata LLL series low ESL ceramic chips (10V). $C1=1\ \mu\text{F}$ and $C2=0.47\ \mu\text{F}$.

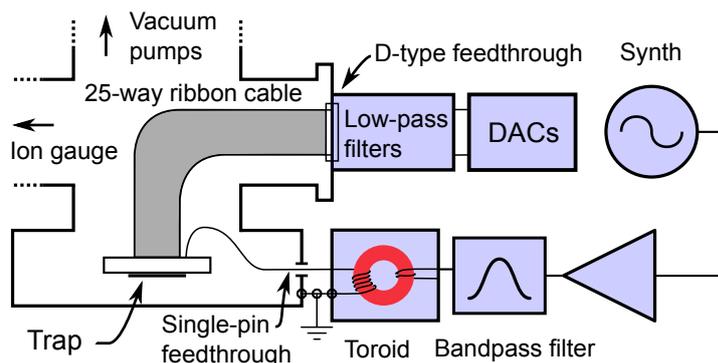


Figure 4.5: Schematic of the vacuum system and wiring arrangement of the trap. The dc electrode voltages were supplied by digital-to-analogue converters (DACs), filtered by low-pass filters, and taken to the trap by a 25-way vacuum-compatible ribbon cable (we were only using the central 10 pairs of electrodes, the rest are grounded on the package). The trap rf drive was provided by a synthesizer, which was amplified and filtered, then stepped-up using a toroid in a resonant transformer arrangement, as described in the text. The vacuum chamber was used as the common grounding point for rf and dc voltages.

noise of $< 100\ \text{nV}/\sqrt{\text{Hz}}$ but this was reduced by a further factor of $> 10^5$ by low-pass filters with cut-off frequency $< 10\ \text{Hz}$ (see fig. 4.4). A surface mount filter board with a continuous ground plane and low-inductance capacitors right up to the feedthrough was found to be much more effective than one fabricated with similar values but using leaded components on Veroboard. The voltages after the filter board were checked for any residual noise using a spectrum analyzer and a custom pre-amplifier with a noise floor of $0.25\ \text{nV}/\sqrt{\text{Hz}}$.

4.3 Operating parameters

As this trap has a similar split central control electrode configuration to the trap used in chapter 3, we use the same techniques to generate voltage sets with

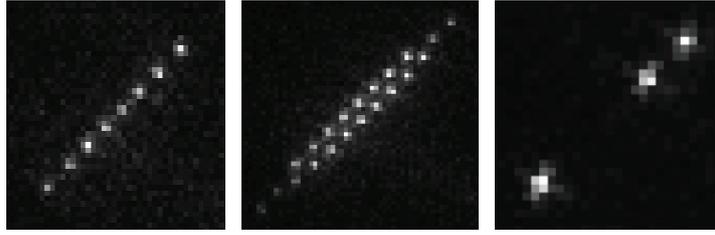


Figure 4.6: 8 ion $^{40}\text{Ca}^+$ crystal (left), 21 ion $^{40}\text{Ca}^+$ crystal (centre) and mixed $^{40}\text{Ca}^+$ and $^{43}\text{Ca}^+$ crystal (right). The $^{43}\text{Ca}^+$ ions are dark because the 397 nm cooling laser was not modulated to span the 3.2 GHz hyperfine splitting of the transition.

the desired axial trapping frequency and radial principal axis orientation. The agreement between the modelled and actual voltages required to obtain a well-compensated trap of given secular frequencies was better than 10% for all electrodes except for the centre dc electrode, for which the measured out of plane field differed by 20% from the model. We attribute this difference to having excluded the dielectric aspects of the trap structure in simulations.

4.4 Ion loading and lifetime

Overall photon detection efficiency was 0.23%. Background scatter was ~ 100 counts/s at this 397 nm power ($2.0 \mu\text{W}$ in a spot with $1/e^2$ radius $w = 30 \mu\text{m}$) [LAS⁺11]. The single-ion trapping lifetime was of order one hour with laser cooling and of order one minute without. In Traps 1 and 2 the two-ion lifetime was only a few minutes (with cooling) and we were unable to load large crystals. In Trap 3 the two-ion lifetime was approximately half the single-ion lifetime and we were able to load large crystals with tens of ions (see fig. 4.6). The difference in behaviour is likely because, with multiple ions, micromotion causes parametric heating when the ions are hot and the motion becomes anharmonic [WMI⁺98b] (for example, after a background gas collision or just after loading). We have also successfully loaded $^{43}\text{Ca}^+$ ions and sympathetically cooled them using $^{40}\text{Ca}^+$ ions [HMS⁺09].

4.5 Micromotion compensation

In Trap 1 we were unable to compensate the rf micromotion with static electric fields, implying that there was no stationary rf null [BMB⁺98]. The residual micromotion was mainly perpendicular to the plane of the trap electrodes, with an amplitude of ~ 200 nm. This was estimated by measuring the amplitude of the micromotion sidebands on a spectrum obtained with a $D_{3/2}-P_{3/2}$ 850 nm repumping beam propagating at 45° to the plane of the trap (see section 3.4). The direction of the motion meant that we could discount a phase difference between the two rf rails, which would cause radial micromotion parallel to the plane of the trap. End effects due to the finite length of the trap were also predicted to be negligible from electric field simulations. This implied that the likely cause of the micromotion was rf pickup on the dc electrodes. For this to be the case the pickup voltages have to be out of phase with the rf electrode voltage, and the secondary rf null created by the pickup has to be in a different place to the principal rf null. We believe that in our case phase shifts were caused by the complex impedance of the ~ 150 mm long ribbon cable (see fig. 4.5) connecting the trap to the vacuum feedthrough or by crosstalk between the wires in the cable.

In an effort to reduce the micromotion in Trap 1 we added a low voltage “compensation” rf voltage to one of the central dc electrodes. This rf was supplied by a synthesizer phase-locked with a variable offset to the main rf drive synth. By adjusting the phase offset and the amplitude of this synthesizer we were able to eliminate the micromotion parallel to the trap surface. This implied that a significant cause of the problem was rf pickup on these centre electrodes, which might be expected since, out of all the dc electrodes, they have by far the largest capacitive coupling to the rf electrodes and therefore will have the largest rf pickup.

On Trap 2, therefore, we added a 1 nF capacitor to the package between each

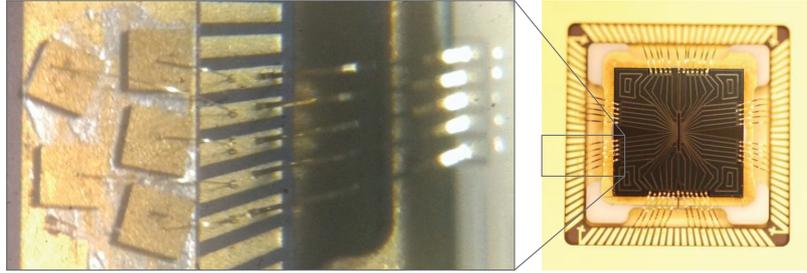


Figure 4.7: Magnified view of capacitors added to the ceramic pin grid array (CPGA). The capacitors were glued to the outer, grounded, gold ring of the CPGA, and then wire bonded to the bond pads as shown.

of these two electrodes and the trap ground plane. However, significant micromotion perpendicular to the trap surface was still present, so on Trap 3 we instead added a 820 pF single-layer capacitor (as in section 3.2.1) between ground and each of the 22 dc electrodes which were in use. These capacitors were glued onto the outer gold ring of the CPGA with Epo-Tek H20E conductive epoxy and then wire bonded to the pads on the CPGA (see fig. 4.7). These capacitors have the effect of moving the common rf grounding point onto the trap package, eliminating differential phase-shifts due to the ribbon cable.

Typical compensation fields were approximately 100 V/m (300 V/m) parallel to (perpendicular to) the plane of the trap, at 2.35 MHz radial secular frequency. The addition of the 820 pF capacitors to Trap 3 reduced the residual micromotion to a similar amplitude to that caused by a 1 V/m (10 V/m) excess field in these respective directions. This implied there was a residual motion of order 1 nm (10 nm), which is at the limit of what we can conveniently resolve with our compensation technique.

4.6 Charging effects

The traps appeared to be fairly resistant to charging under normal operation. The field compensation in the direction parallel to the trap plane drifted by only ~ 1 V/m on an hour timescale and we observed no change in the field compen-

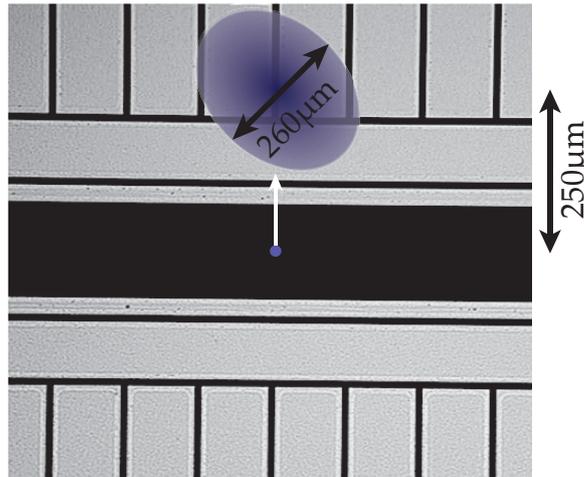


Figure 4.8: Diagram to show how the 397 nm charging beam strikes the trap. The beam has a circular profile with spot size $w = 130 \mu\text{m}$, but propagates at 45° to the surface so it illuminates an elliptical area of the electrodes. The white arrow shows the direction of the induced electric field measured at the ion (blue dot).

sation during or after ion re-loading.

A more detailed study of the trap charging was conducted by irradiating the electrodes directly with another 397 nm laser beam to induce a larger charging effect. The beam had a power of $10 \mu\text{W}$ in a circular spot with $1/e^2$ radius $w = 130 \mu\text{m}$ and propagated at 45° to the surface (see fig. 4.8). Data were taken by turning the laser on for a fixed amount of time, then repeatedly recompensating the in-plane micromotion at ~ 30 s intervals. This method has a resolution of ± 1 V/m with 10 s detection time. Both the direction of the resulting field (which attracted the ion towards the beam spot) and the time dependence of its decay were in qualitative agreement with previous work by Harlander and coworkers, for a surface trap with copper electrodes [HBHB10]. Our data are well fitted by a double exponential (see fig. 4.9) with time constants 78 and 650 seconds compared with 5 and 120 seconds reported in that work. The slower discharge time may arise because the native oxide layer on our aluminium trap was thicker or less conductive than that on the copper trap in [HBHB10] and could support a longer relaxation time for any charges on it [RMG77].

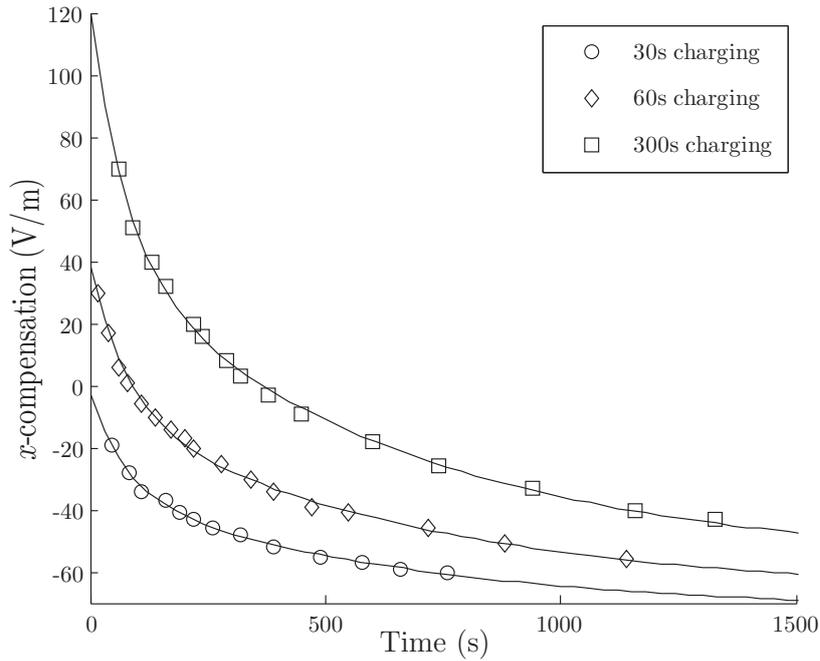


Figure 4.9: Electric field measured at the ion, after irradiation of the trap by 30, 60 and 300 s exposure to the 397 nm charging beam. The curves are a joint fit to $E_x = A_i + B_i e^{-t/t_1} + C_i e^{-t/t_2}$ and the fitted time constants are $t_1 = 78(20)$ s and $t_2 = 650(80)$ s. The errors on each measurement are approximately the size of the symbols used.

4.7 Heating rates

We made several studies of the heating rate in Trap 3 using our simplified method for Λ systems which is described in section 3.5. This uses a repumping scheme which reduces the Λ system to an approximate two-level atom such that the standard one-dimensional Doppler-recooling model of [WEL⁺07] is applicable. The Doppler cooling beam is sensitive to heating along all three principal axes and hence cannot distinguish axial and radial heating, but if the axial frequency is significantly below the radial frequencies, the dominant contribution to the signal is expected to be from axial heating as the electric field noise has generally been found to scale approximately as $1/f$. We tested this by measuring the heating rate at various radial frequencies (fig. 4.10a); we observe no significant change in the heating rate; this suggests that the heating is dominated by the axial contribution. We further measure no increase in the heating rate with

variations in the out-of-plane compensation field of up to ~ 100 V/m.

The heating rate dependence on the axial frequency (see fig. 4.10b) is comparable with the $1/f$ dependence reported in the literature [DOS⁺06, ESL⁺07, LGA⁺08, TKK⁺00]. However the poor fit to a $1/f^n$ curve, especially at low frequencies, implies the noise in our system does not have a simple $1/f^n$ dependence. See chapter 5 for further discussion of frequency dependencies.

The data in fig. 4.10 were taken at the centre of the trap but we also loaded ions at locations ± 270 μm along the trap z -axis and measured the same heating rates there. We observed no change in the heating rate over a period of several months, indicating that ion loading did not adversely affect the electric field noise (in contrast for example with the observations of [DNM⁺11]).

We also attempted to measure the heating rate in Trap 2 but it exhibited a strong dependence on the rf drive amplitude pointing to a heating contribution from the radial modes (in which case the simple one-dimensional heating rate model is inapplicable). The uncompensated radial micromotion in Trap 2 may be responsible for this because an uncompensated trap is very sensitive to noise at $\Omega_{\text{rf}} \pm \omega_{\text{r}}$ [Bla10]. This is supported by the observation that we found a large increase in heating rate when the rf filter was removed. In contrast, in Trap 3 we measured no difference in heating rate with or without the rf filter.

4.8 Future improvements

4.8.1 Coated electrodes

While the charging effects measured in section 4.6 are negligible for light scattered from low-power (few μW) Doppler-cooling or optical pumping beams, they could be an issue for ~ 100 mW beams potentially required for implementing quantum logic gates [OIB⁺07]. Comparisons made by Wang and coworkers [WLL⁺11] indicate that coating the trap with another material, such as gold,

4. CHARACTERISATION OF A MICROFABRICATED TRAP

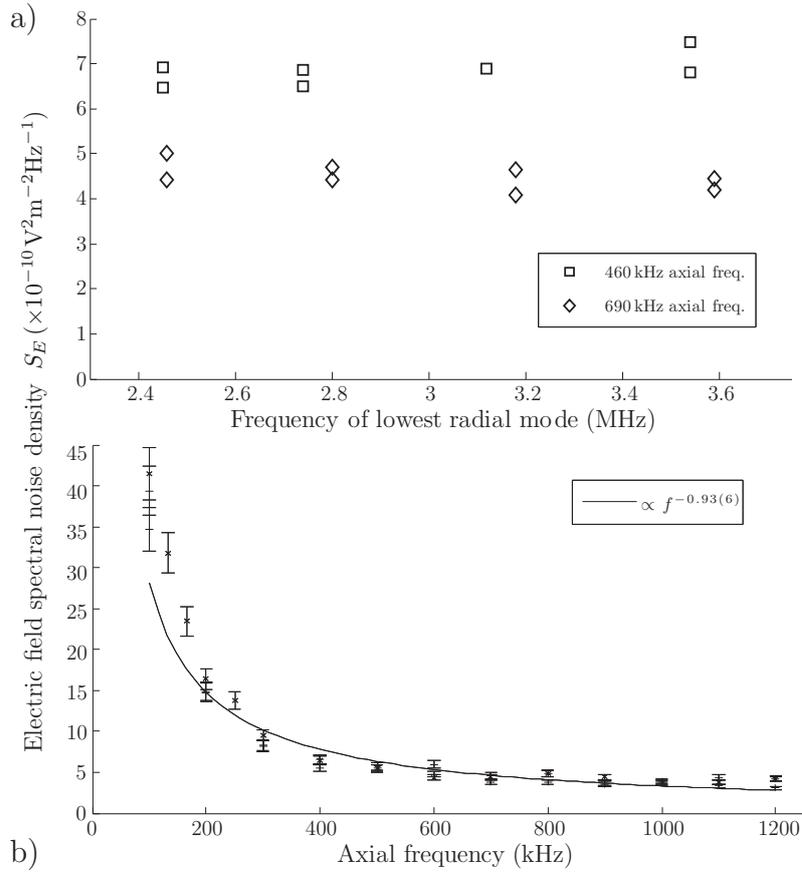


Figure 4.10: Single-ion heating rate in Trap 3, expressed as the electric field spectral noise density experienced by the ion. (a) Heating rate versus radial frequency at fixed axial frequencies of $f_z = 460$ kHz and $f_z = 690$ kHz. (b) Heating rate versus axial frequency at fixed rf amplitude ($f_{r1} \simeq 3.1$ MHz, $f_{r2} \simeq 3.5$ MHz). The data points marked (\times) were taken a week after those marked ($+$). The curve is the weighted least-squares fit to a $1/f^n$ power law, giving $n = 0.93(6)$. The heating rate equates to ~ 55 quanta/ms at 1 MHz axial frequency.

which does not support a native oxide layer, will reduce the charging measured in section 4.6 by more than an order of magnitude. For this design of trap, evaporative coating is straightforward as the electrodes themselves act as a shadow mask to prevent shorting. As this evaporation can be done as a final processing step after packaging, even materials like gold (see fig. 4.11) which are not compatible with CMOS fabrication can be used. A similar gold coated trap has been tested at Sandia and ions have been loaded into it and shuttled without problems, though detailed studies are yet to be carried out.

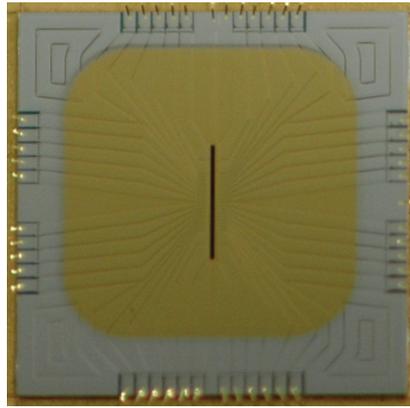


Figure 4.11: A similar trap where the ion trapping region has been evaporatively coated with gold by Sandia.

4.8.2 Integrated filters

As demonstrated in section 4.5, it is necessary to minimize rf pickup on control electrodes, and hence placement of capacitive filters in the closest possible proximity to the electrodes is optimal. Additionally, filtering with a cut-off frequency larger than the DAC update rate and smaller than the rf drive and ion secular frequencies is required. For a 1 MHz cut-off frequency, a capacitor value of 1 nF and a resistor value of 1 k Ω can be integrated directly into the trap chip itself as part of the chip fabrication process. Sandia have successfully fabricated “trench” capacitors (see fig. 4.12) with a capacitance of 94 fF/ μm^2 which can be integrated into the next iteration of these traps [AHJ⁺11].

4.9 Comparison with other work

The heating rate observed in this trap is comparable to that measured in other traps of similar ion-electrode distance scale [ABLW11]. In comparison with the microfabricated aluminium-electrode trap reported in [LLC⁺09], the present trap has a heating rate approximately a factor of 5 times worse at 1 MHz, but we demonstrate a much longer trapping lifetime, reduced stray fields and have experienced no trap reliability issues (see also [Cur10] for details of an unsuc-

4. CHARACTERISATION OF A MICROFABRICATED TRAP

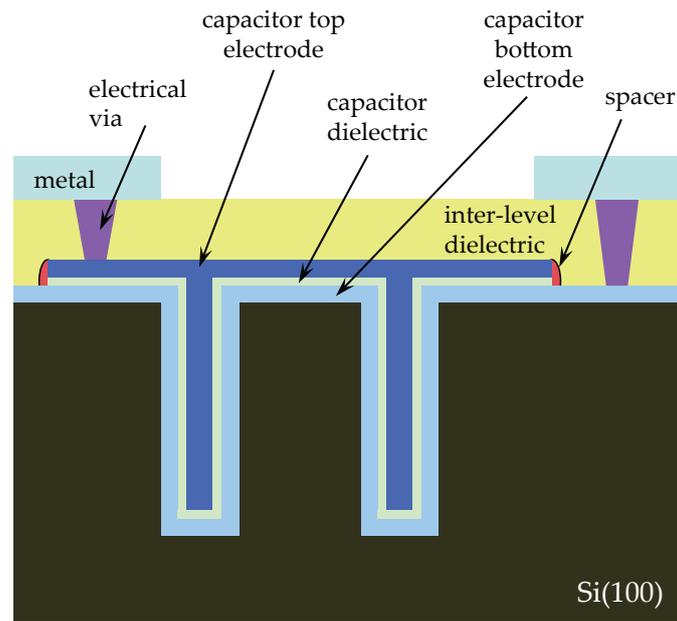


Figure 4.12: Schematic of trench capacitors built into the ion trap chip Si substrate.

cessful attempt to use this trap by our group).

5

Pulsed laser cleaning

I was talking to a guy the other day who was trying to convince me that CDs were better than vinyl because they had no surface noise. And I said “listen mate, life has surface noise.”

John Peel

5.1 Anomalous heating in ion traps

As discussed in chapter 1, the presence of uncontrolled fluctuating electric fields in ion traps affects the ions’ external motion and induces an “anomalous heating” that limits the achievable fidelity of multi-ion quantum gates that rely on the coherent control of this motion [WMI⁺98b]. Experimental observations concerning this phenomenon are consistent with a very unfavorable d^{-4} scaling law with distance from the substrate, compatible with a random distribution of fluctuating charges or dipolar “patches” at the electrode surfaces [TKK⁺00, DHL⁺04]. In addition, the scaling of the electric field noise spectral density $S_E(\omega)$ with respect to trap secular frequency ω has been found to approximate a $\omega^{-\alpha}$ law, with exponents roughly compatible with $\alpha = 1$ but spanning the range $0.4 < \alpha < 1.6$ (see section 4.7 and [TKK⁺00, DHL⁺04, DOS⁺06, ESL⁺07, LGL⁺08]).

5.1.1 Theoretical models

Some recent theoretical models propose the fluctuations of the electric dipoles of adsorbed molecules as a possible driving mechanism [DNM⁺11, SNRWS11], whilst other authors point out the role played by a more general (but microscopically not identified) correlation length associated with disorder on the surface [DCG09, LHC11]. Several studies point to surface contamination being an issue: the first observed an order of magnitude variation between four nominally identical traps and even the same trap after it had been re-cleaned [LGL⁺08], the second observed no change in the heating rate even when the bulk of the electrode undergoes a transition to a superconducting state [WGL⁺10], and the third reported an increased heating rate in the region of the trap used for loading ions [DNM⁺11].

5.1.2 Cleaning methods

Although traps are typically cleaned using some combination of organic solvents, ozone cleaning or plasma cleaning after fabrication, any surface that has been exposed to atmosphere will have a covering of adsorbents at least several mono-layers thick. Additionally, trap electrode materials that react with oxygen will have a native oxide layer. Standard methods of preparing atomically clean metal surfaces under ultra-high vacuum involve either *in situ* cleaving, evaporation, or repeated cycles of ion bombardment and high temperature annealing [WD94]. Whilst the latter two processes could be used in principle for micro-traps, they would add significant engineering complexity to the trap structures and to the experimental vacuum systems. Moreover, such harsh treatments are unlikely to be compatible with traps currently under development which incorporate integrated optics [SNJ⁺11, BEM⁺11], an important step towards scaling up ion trap quantum computing.

The cleaning of metallic surfaces based on pulsed-laser sources has also been

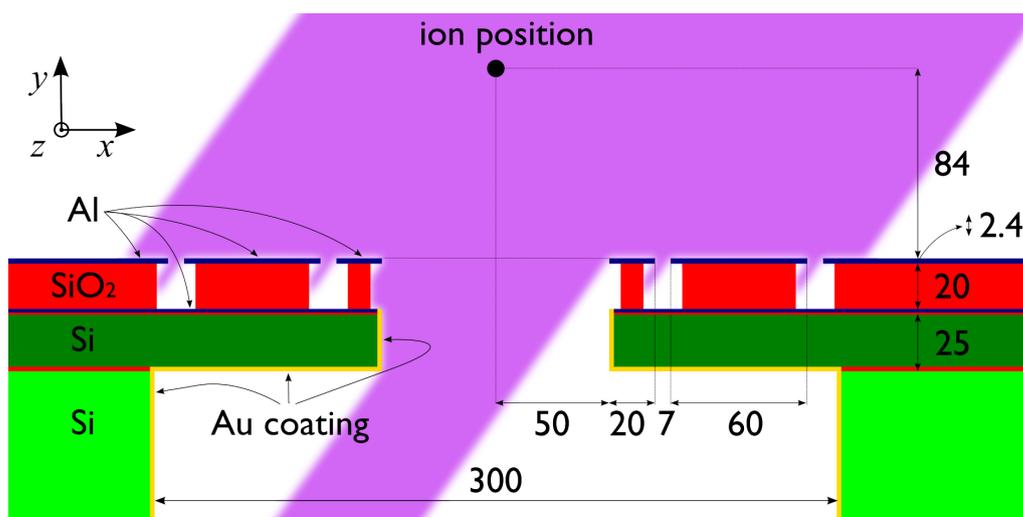


Figure 5.1: Schematic cross-section (to scale) of the microfabricated trap (see chapter 4 and [SFH⁺10]); all dimensions are in microns. The slot through the centre of the substrate, parallel to the trap z axis, allows ions to be loaded from the underside of the substrate: this avoids contamination of the trap electrodes from the neutral atomic beam source. An important issue for the current study is the gold coating that covers the silicon, which has a nominal thickness of $0.11 \mu\text{m}$ on the slot side walls (not to scale on the scheme). The cleaning beam cross-section corresponding to the $\vec{k}_- = (-\frac{1}{2}, -\frac{1}{\sqrt{2}}, \frac{1}{2})$ propagation direction (violet shade) is also sketched.

noted as an effective, if less frequently used, technique for producing clean surfaces [DO06]. The technique is based on the fact that energy density thresholds for desorbing contaminants or removing oxides are generally lower than the ablation damage threshold for the metallic surface [Phi07]. In particular, “dry” laser cleaning, compatible with ultra-high vacuum techniques, has been used for oxide removal from metallic surfaces [MADA99, LW00] and cleaning of aluminium-coated optical surfaces [MWRM96]. Typically, ultraviolet pulses from nanosecond sources (e.g. excimer or Nd:YAG 3rd/4th harmonic) and energy densities of $\sim 100 \text{ mJ}/\text{cm}^2$ are used. Laser-cleaning may be easily applied to an ion microtrap and it requires no modification to a typical vacuum system (viewports which transmit UV are often required for laser access). Furthermore, the cleaning-laser beam can be positioned with micron-level precision and its direction can be easily adjusted, in order to avoid delicate components or to irradiate selectively different parts of complex 3D trap designs.

5.2 Experimental methods

We implement this technique on a state-of-the-art microfabricated trap (see chapter 4 and [BEM⁺11, SFH⁺10]). The structure of the trap (see figure 5.1) is such that three different materials are exposed to the cleaning beam: the aluminium of the upper electrode surface (2.4 μm of sputter-deposited Al-1/2% Cu with 2–3 nm native oxide having an rms surface roughness of ~ 8 nm), the gold coating on the silicon (500 nm Au / 50 nm Pt / 20 nm Ti stack, e-beam evaporated) and the silicon dioxide¹ of the pillars which support the electrodes. We note that the gold coating was evaporated at an angle such that it has a nominal thickness of 114 nm on the slot side walls.

5.2.1 Expected cleaning thresholds

Previous works describe laser ablation or laser cleaning of such materials; we briefly review here the main results that may apply to the present study. For aluminium the generally accepted ablation threshold (for plasma generation) lies around 4 J/cm^2 at $\lambda = 355 \text{ nm}$ [PRM96]; however, careful studies in high vacuum demonstrated that a measurable Al^+ ion yield appears at energy densities lower than 100 mJ/cm^2 [TH00]. At the same wavelength the reported thresholds for cleaning and damaging the surface of an Al coated glass substrate (BK7) are 200 mJ/cm^2 and 490 mJ/cm^2 respectively [MWRM96]. The laser ablation of aluminium oxide is somewhat more complicated due to the fact that several phases can coexist in native oxides. An experiment performed in ultra-high vacuum conditions (ion detection and surface analysis) on a sapphire monocrystal [SH92] gives an ablation threshold of 3 J/cm^2 and a threshold more than one order of magnitude lower for Al^+ ion emission. Even lower thresholds are expected for native oxides [MADA99]. In the case of gold, an experimental study addressed the case of thin films (up to some microns) in

¹plasma deposited TEOS (tetraethyl orthosilicate)

the single-shot regime [MRS⁺94] and estimated the ablation threshold to be $\sim 250 \text{ mJ/cm}^2$ (with the damage threshold a factor of 2 below this) for a film thickness of 100 nm. While the ablation threshold for silicon is well-documented (1.3 J/cm^2) [GKH95], the case of silicon dioxide is less straightforward to analyze, due to the differences in composition and porosity.

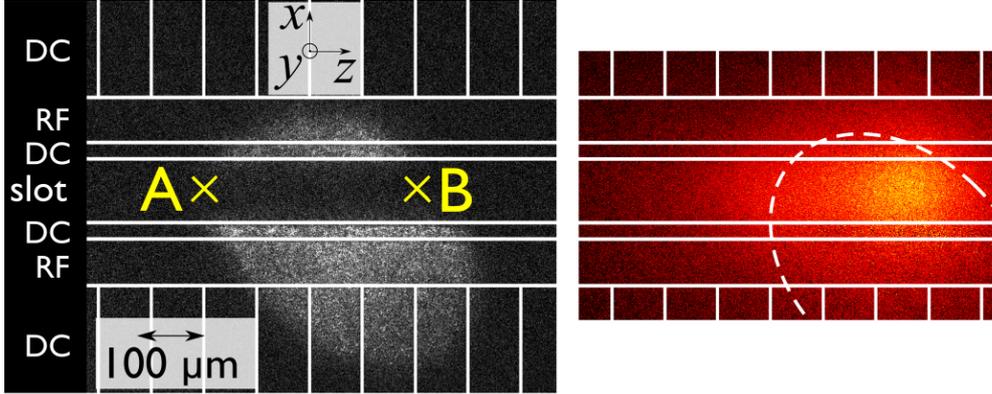


Figure 5.2: Left: typical image of the cleaning laser light scattered from the trap. The electrode gaps are sketched as lines and the two trapping sites “A” and “B” discussed in the text are indicated. Right: plume fluorescence (false colours) associated with a single cleaning pulse (energy density $\simeq 200 \text{ mJ/cm}^2$) impinging on an uncleaned area. The displayed image does not show the scattered laser light which was subtracted using an image taken after cleaning.

5.2.2 Cleaning procedure

The cleaning beam is generated by a tripled Nd:YAG laser² that delivers 3–5 ns pulses (nominal) at $\lambda = 355 \text{ nm}$ with an energy continuously adjustable up to $\sim 1 \text{ mJ}$ and a repetition rate of up to 15 Hz. The beam is spectrally filtered by a fused silica prism and then sent to the trap. A $\sim 300 \mu\text{m}$ diameter pin-hole selects the central part of the laser beam and is imaged on the trap plane in order to obtain a well-defined spot with an intensity which is uniform to $\sim 20\%$. The imaging lens ($2f$ – $2f$ configuration) is mounted on a micrometer translation stage to allow for fine positioning of the cleaning spot. In view of the particular geometry of the trap, with a slot through the centre of the substrate,

²Continuum Minilite ML I

two symmetric beam paths, both at 45° incidence to the substrate plane (xz -plane), are used to allow cleaning of both interior walls of the slot (see figure 5.1). The spot size and position on the trap are monitored using an EM-CCD camera (see section 2.3.2), also used for imaging of the trapped ion. A typical image of the UV light scattered from the trap is shown in figure 5.2 where the geometry of the electrodes is also sketched. With this setup, energy densities up to $\simeq 350 \text{ mJ/cm}^2$ can be obtained with a spot diameter of $\sim 300 \mu\text{m}$. Here and in the following, energy densities are given normal to the beam propagation direction. The energy density on a specific trap surface is reduced by a geometrical factor: for the xz -plane (upper electrode surfaces) this factor is $\frac{1}{\sqrt{2}}$, for the yz -plane (slot side walls) it is $\frac{1}{2}$. All trap electrodes were grounded while firing the laser, to prevent the possibility of arcing initiated by photoelectrons.

5.3 Results

5.3.1 Ion heating rate study

Before applying any laser cleaning, we characterized the trap heating rate by testing three trapping sites ($z = 0, \pm 240 \mu\text{m}$ from centre). The heating rate and frequency dependence were uniform (within the estimated error) and compatible with the previous measurements on the same trap given in section 4.7. Contrary to the case of [DNM⁺11], we did not observe an increase of the heating rate over an operation time of several months. In our loading geometry the oven is placed below the slot shown in figure 5.1, $\sim 50 \text{ mm}$ behind the trap. Based on data from a similar oven [LRH⁺04], we can estimate the order of magnitude of the flux reaching the interior walls of the slot when the oven is on: $\sim 10^5 \text{ Ca}$ atoms per second and per mm^2 (less than 1 monolayer every 2 years). However, we can not exclude a contamination of the slot surfaces by other species during the initial firing of the oven: at that moment, while the Ca flux was still $\lesssim 10^7$

(atoms/s)/mm², the pressure in the vacuum chamber increased up to 10⁻⁸ mbar (see section 2.5.1).

We began the laser cleaning by applying pulses to a trapping region two electrodes away from the centre of the trap array ($z = +160 \mu\text{m}$, site “B” on figure 5.2). Each experiment consisted of applying a number of pulses at a given energy density (1 Hz repetition rate) around the trapping position. After each experiment the trap was loaded and the ion’s heating rate was measured. Then the energy density was increased for the next experiment. At 30 mJ/cm² the DC electrode centred at $x = -60 \mu\text{m}$ showed clear signs of delamination near $z = +400 \mu\text{m}$, presumably caused by differential expansion induced by heating (see figure 5.5a). At this point no change in heating rates had been observed. As we did not want to risk further damage to the trap we moved the trapping region to a symmetric position 4 electrodes away (site “A”, $z = -160 \mu\text{m}$) and resumed the experiment, reducing the repetition rate to 0.2 Hz to minimize the risk of heat accumulation. As we were still able to trap in site “B” and the heating rate had not significantly changed there, we later used it as a control measurement to ensure that any measured change in heating rate at zone “A” over time was not due to a systematic effect in our measurement or a change in some global noise source (e.g. electrical pickup).

Figure 5.3 shows the evolution of the heating rates (expressed in terms of electric field noise spectral density) and micromotion compensation fields throughout the entire series of cleaning experiments applied to site “A”. Initially we applied the cleaning beam along the $\vec{k}_+ = (\frac{1}{2}, -\frac{1}{\sqrt{2}}, -\frac{1}{2})$ direction only (see figures 5.1 and 5.2), indicated by black bars in figure 5.3. There appears to be a slight drop in heating rates from $\sim 100 \text{ mJ/cm}^2$ onwards, which initially is not much below the scatter on the measurements. However, once we attempted cleaning also from the $\vec{k}_- = (-\frac{1}{2}, -\frac{1}{\sqrt{2}}, \frac{1}{2})$ direction (indicated by red bars in figure 5.3) with an energy density of 100 mJ/cm², the drop became much

more pronounced. This effect points to a large contribution to the noise from the slot side wall, the only significant area not cleaned by the \vec{k}_+ directed beam. When cleaning a “fresh” region, we also observed for each single-pulse shot a fluorescent emission (ablation plume) from inside the slot (see figure 5.2) and an accompanying pressure spike of a few 10^{-12} Torr. The plume fluorescence intensity and the pressure spike amplitude dropped rapidly and became undetectable after 3 or 4 shots, implying the source material responsible for these phenomena had been removed. These effects were not observed in the first \vec{k}_+ cleaning direction. This is due to the fact that along the \vec{k}_+ direction we gradually increased the intensity over thousands of pulses: this probably removed the material in smaller amounts, below the sensitivity of the camera or ion gauge.

At this point, heating rate data as a function of axial trap frequency was taken at both site “A” and “B” (see figure 5.4). Whilst the heating rate in site “B” is still entirely consistent with the data taken several months before (section 4.7), that in site “A” shows a marked decrease and also a significant drop of the exponent α .

The exposure of the trap to cleaning laser pulses also caused a shift in the micromotion compensation voltages along both x and y directions. These shifts had a small component ($\sim 10\%$) which relaxed over several hours (presumably induced by charging (see section 4.6 and [HBHB10]) but the major part of the effect did not relax, even over weeks. The direction of the electric field to be compensated was mainly such that the ion was attracted upwards ($+y$) and away from the side of the slot being cleaned. The effect appeared to have somewhat saturated until the \vec{k}_- cleaning direction was used, at which point the field roughly doubled in magnitude along y but evened out in x . Again, this behaviour points to a major contribution from the large slot side wall (silica or gold surfaces).

We then attempted to reduce the heating rate further by increasing the en-

ergy density to 360 mJ/cm^2 in the \vec{k}_- direction; however this caused visible damage to the aluminium top surface of the trap (observed as an increase of light scattering in the irradiated zone). This damage caused an increase in the heating rate (though still below the initial value) and a reversion to a higher exponent α in the frequency dependence [$\alpha = 0.88(3)$].

5.3.2 Electron microscopy study

The still-operational microtrap was then removed from the vacuum chamber and observed under optical and electron microscopes. Optical microscope images confirmed some visual damage of the Al surface of the electrodes surrounding site “A” and suggested a reflectivity decrease of the slot side wall where it had been irradiated. Electron microscope images were taken both with secondary electron and back-scattered electron (BSE) contrasts. As shown in figure 5.5 a, the delamination of the DC electrode (at $z \sim +400 \text{ }\mu\text{m}$) appears to be associated with a delamination of the silica pillar, suggesting that some thermally-induced stress may be at the origin of this damage. The boundary between irradiated (but not damaged) and non-irradiated region inside the slot side wall displays some change in the topography of the gold coating (figure 5.5 b). However, an image obtained with BSE contrast (which is sensitive to the atomic number, Z) shows that the gold was probably only removed (in a stripe-like fashion) around the damaged site “A” (figure 5.5 c) and it still forms a continuous film in the regions irradiated with $< 200 \text{ mJ/cm}^2$ energy density.

5.4 Discussion

This study shows that the technique of high-intensity laser irradiation is capable of reducing *in situ* the heating rate of a microfabricated ion trap. It is also notable that the exponent α , characterizing the electric field noise frequency dependence, is affected by the procedure (figure 5.4).

Two possible interpretations of the mechanism involved in this electric field noise reduction can be pointed out. The first one is based on the (possibly partial) removal of surface contamination, responsible for the existence of patches [DNM⁺11, SNRWS11]. The theoretical study in [SNRWS11] suggests that different adsorbates could give rise to different frequency dependences: partial cleaning of a sub-set of adsorbates could then explain the observed change in the exponent α . The second interpretation is that the observed effect was caused by the apparent change in the topography of the thin gold film inside the slot side wall (see figure 5.5 b). According to [DCG09, LHC11], this re-arrangement of the metallic film could increase the characteristic length ζ of the disorder, leading to a reduction of the heating-rate. It should also be noted that in spite of the laser-cleaning procedure, we were unable to bring the measured heating rate below the best results obtained with traps of this size at room temperature ($S_E \sim 2 \times 10^{-12} \text{ V}^2\text{m}^{-2}\text{Hz}^{-1}$ at $\omega/2\pi = 1 \text{ MHz}$, cf. figure 5 of [DNM⁺11]). It is likely that surface contamination is only one contributing factor to anomalous heating. If so, this technique could still be very useful as a method of reducing the large variance observed between traps of the same material and fabrication which currently renders any systematic studies into the best material and fabrication choice very difficult.

The particular microtrap that we used for this investigation was not ideally suited to the purpose; different materials were irradiated at the same time and the ion still had a direct line-of-sight to the dielectric pillars. An improved version of the trap with shorter dielectric pillars and a front, as well as back, evaporated coating has already been demonstrated [AHJ⁺11]. A gold coating was used in that case but in principle any conducting material could be used. A similar experiment with such a trap would be easier to interpret as only this one material predominates.

In order to develop this technique further a deeper understanding of clean-

ing and damage thresholds for typical trap structures and adsorbates is needed. This could be done by combining laser-cleaning with analysis of the surface chemical composition (e.g. by Auger or X-ray photoelectron spectroscopy techniques). Lasers with a higher photon energy (e.g. Nd:YAG 4th harmonic or excimer) or better ratio between peak intensity and average power (e.g. femtosecond lasers) should be investigated because more effective cleaning is expected for an equivalent thermal load. Following these lines, an optimal combination of electrode materials and cleaning laser could be identified, allowing for routine *in situ* cleaning of microtraps whenever necessary.

5. PULSED LASER CLEANING

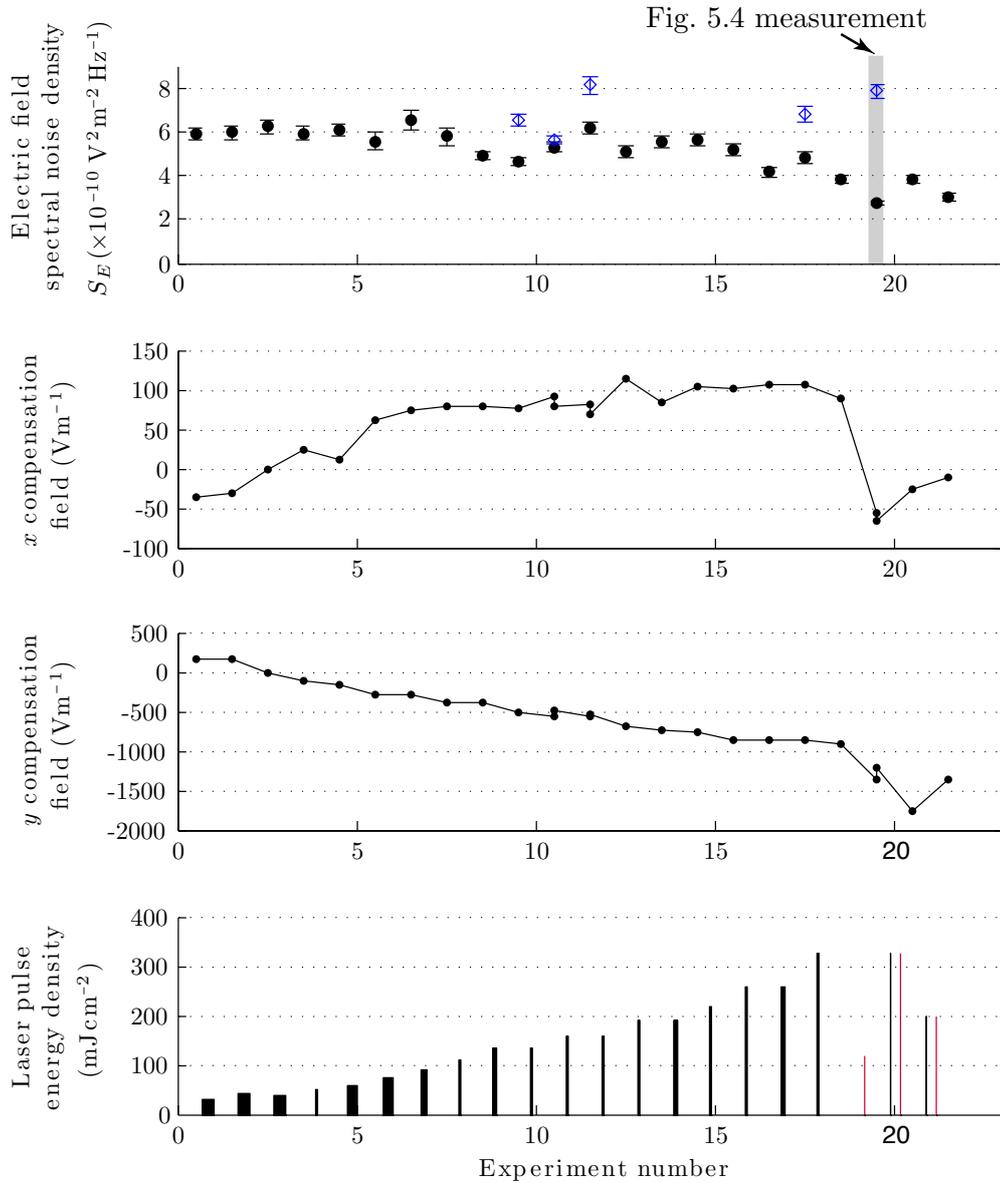


Figure 5.3: Top to bottom: heating rates for an axial trapping frequency $\omega_z/2\pi = 500$ kHz (expressed in terms of electric field noise spectral density), micromotion compensation fields (x and y directions) and cleaning laser energy density plotted against the experiment number. Each experiment consisted of a cleaning attempt followed by micromotion compensation and heating-rate measurements. In the top graph, blue open diamonds correspond to control measurements in site “B” while black filled circles correspond to measurements in the cleaned site “A”. Error bars are derived from the scatter in measurements repeated under nominally identical conditions. In the bottom graph, the number of cleaning pulses applied in each experiment is proportional to the thickness of the bar (1000, 400 or 100 pulses); black (red) bars indicate cleaning from the k_+ (k_-) direction, respectively (see text). The entire data set was taken over a ten week period, and the repetition rate was 0.2 Hz throughout. The pulse energy density is given normal to the cleaning beam direction.

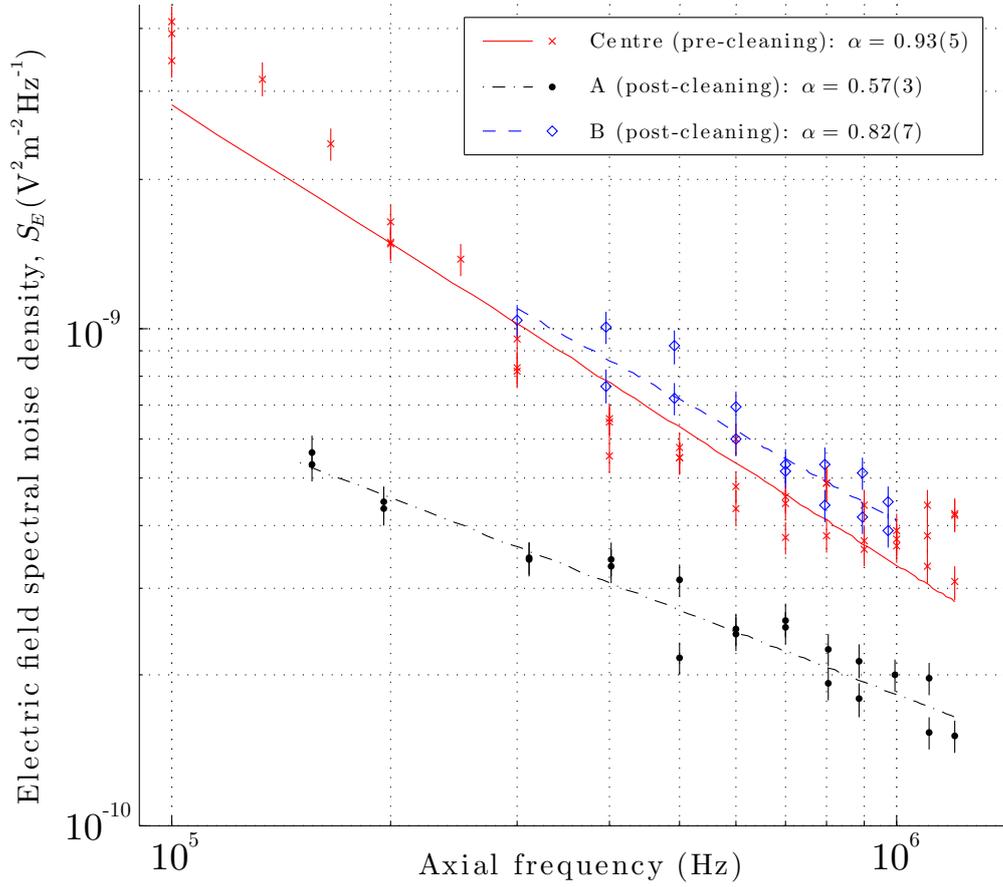


Figure 5.4: Heating rate (expressed in terms of electric field noise spectral density) as a function of the axial trapping frequency $\omega_z/2\pi$. We compare two data sets obtained in the “A” and “B” trapping sites (black filled circles and blue open diamonds, respectively) taken on the same day with the same settings. For reference, the data corresponding to the central site of the trap taken before cleaning in section 4.7 (red crosses) are also shown. Error bars are derived from the scatter in repeated measurement sets and the lines correspond to the best fits of each data set to a $\omega_z^{-\alpha}$ law. The exponents α corresponding to the “A”, “B” and “centre” sites are 0.57(3), 0.82(7) and 0.93(5) respectively.

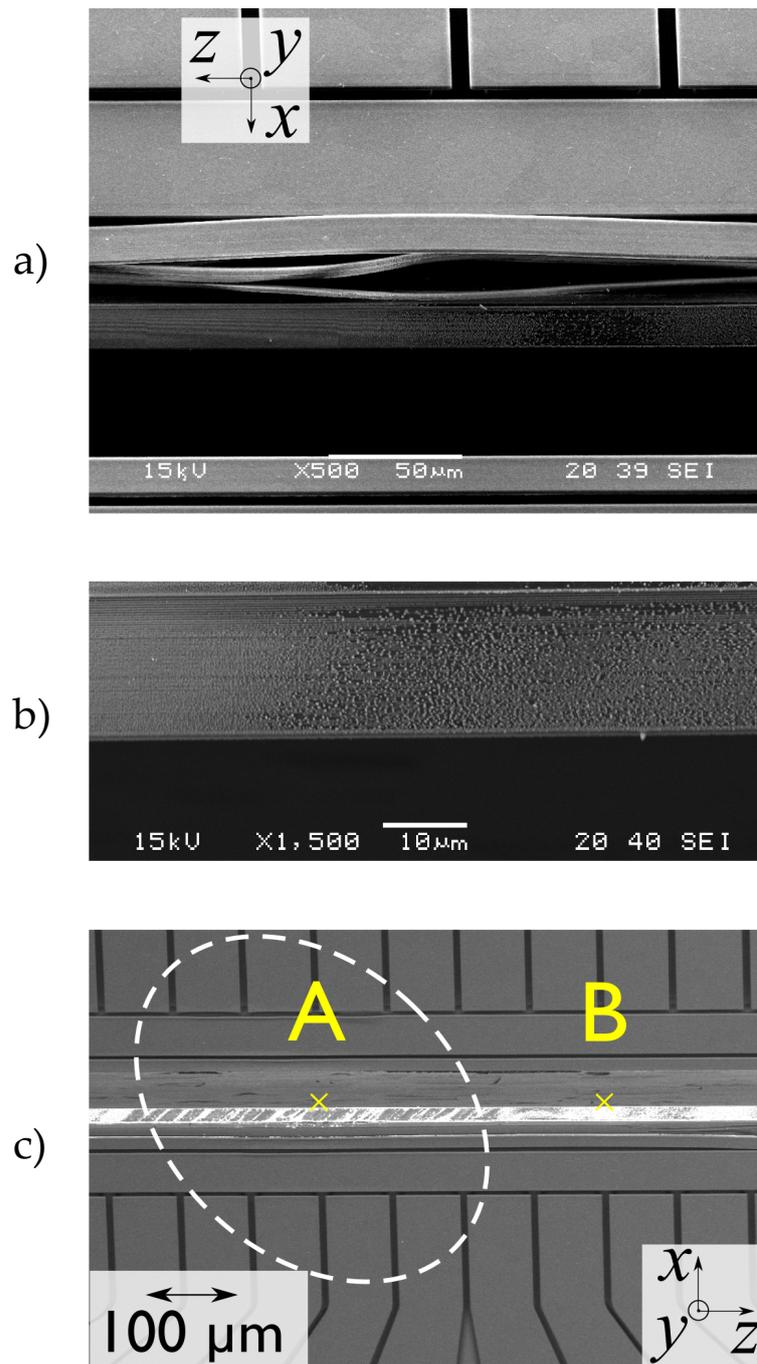


Figure 5.5: Scanning electron microscope images of the microtrap (under a 45° angle to the y -axis, giving a view of a side wall of the slot) obtained after laser-cleaning experiments. a): delamination damage of DC electrode at $z = +400 \mu\text{m}$. b): topography change between non-irradiated (left) and irradiated (right) zones of the side wall. c): BSE contrast image of the side wall around the trapping site "A". This side wall was irradiated with the maximum nominal energy density of $\approx 360 \text{ mJ}/\text{cm}^2$, i.e. $\approx 180 \text{ mJ}/\text{cm}^2$ after the geometrical correction. Due to the Z -contrast of the BSE the gold shows up bright; it appears to have been completely removed in places. (the corresponding laser spot is sketched for reference)

6

Qubits for microwave-driven quantum logic in $^{43}\text{Ca}^+$

Rip it up and start again.

Orange Juice

A scheme to drive multi-qubit entanglement in ions was with near-field microwaves was proposed in 2008 by Ospelkaus et al [OLA⁺08] and has recently been experimentally demonstrated in $^{25}\text{Mg}^+$ [OWC⁺11]. In this chapter we consider possible qubits that could be used to implement the same scheme in $^{43}\text{Ca}^+$.

6.1 Hyperfine qubits

In order to carry out quantum logic driven by an oscillating magnetic field we require a qubit transition at radio or microwave frequency. Hyperfine or Zeeman states in the ground level of an atom are in this regime and make good qubits because they are extremely long lived. Hyperfine qubits have a further advantage that magnetic-field-independent transitions, known as ‘clock’ transitions, exist. Qubits based on these transitions can have coherence times of over ten seconds [LOJ⁺05]. In order to have a ground level hyperfine qubit an ion must have a nuclear magnetic moment and the ground level must have $J > 0$. Table 6.1 gives a

list of all the experimentally trapped and laser cooled ions that meet this criteria. In all of them, the $M = 0$ to $M = 0$ transition is field-independent at zero field. For ions with $I > 1/2$ there are other transitions that are field-independent at a non-zero field. These ‘intermediate-field’ clock qubits have been successfully demonstrated in $^9\text{Be}^+$ [LOJ⁺05] and $^{25}\text{Mg}^+$ [OWC⁺11] at NIST.

Table 6.1: A summary of ground level hyperfine qubits in experimentally trapped ions and example references. Some ions have not been used for QIP and the reference refers to metrology or spectroscopy experiments. All the ions listed have a $^2S_{1/2}$ ground level.

Ion	f_{qubit} (GHz)	I	Cooling wavelengths (nm)	Reference
$^9\text{Be}^+$	1.25	3/2	313	[MMK ⁺ 95]
$^{25}\text{Mg}^+$	1.69	5/2	279	[FSG ⁺ 08]
$^{43}\text{Ca}^+$	3.23	7/2	397, 866	[LKH ⁺ 07]
$^{87}\text{Sr}^+$	5.00	9/2	422, 1092	[BGG ⁺ 03]
$^{111,113}\text{Cd}^+$	14.5, 15.3	1/2	214	[HLB ⁺ 05]
$^{135,137}\text{Ba}^+$	7.18, 8.04	3/2	493, 650	[DKN ⁺ 10]
$^{171}\text{Yb}^+$	12.6	1/2	370, 935	[OYM ⁺ 07]
$^{173}\text{Yb}^+$	10.5	5/2	370, 935	[MBG ⁺ 87]
$^{199}\text{Hg}^+$	40.5	1/2	194	[BMC ⁺ 98]
$^{201}\text{Hg}^+$	30.0	3/2	194	[BTLT09]

6.2 The $^2S_{1/2}$ ground level of $^{43}\text{Ca}^+$

6.2.1 Structure

The qubit transitions we will consider in this chapter exist in the intermediate regime where the interaction energy with the external magnetic field is neither weak nor strong compared with the hyperfine interaction. This means that the states have to be found as linear combinations of some set of basis states. For later convenience we will work in the M_I, M_J basis. The projection of the angular momentum onto the external field M remains a good quantum number at all fields and $J = 1/2$ so any state can be a superposition of at most two basis

states:

$$\begin{aligned}
 |n(M_n)\rangle &= \alpha_n |JI, M_J = -\frac{1}{2}, M_I = M_n + \frac{1}{2}\rangle \\
 &\quad + \beta_n |JI, M_J = \frac{1}{2}, M_I = M_n - \frac{1}{2}\rangle \\
 &\equiv \alpha_n |-\frac{1}{2}, M_n + \frac{1}{2}\rangle + \beta_n |\frac{1}{2}, M_n - \frac{1}{2}\rangle.
 \end{aligned} \tag{6.1}$$

There are two eigenstates for each value of M and we label the higher energy states n (as above) and the lower energy states m . The exception are the $M = \pm 4$ states which are the pure states $|\pm 1/2, \pm 7/2\rangle$. The Breit-Rabi solution gives a simple analytical expression for the energies of the eigenstates

$$E(B_0) = \frac{-E_{hfs}}{2(2I+1)} - g_I \mu_N B_0 M \pm \frac{E_{hfs}}{2} \sqrt{1 + \frac{4\chi B_0 M}{2I+1} + \chi^2 B_0^2} \tag{6.2}$$

where $\chi = (g_I \mu_N + g_J \mu_B)/E_{hfs}$. μ_B is the Bohr magneton, μ_N is the nuclear magneton, M is the magnetic quantum number and B_0 is the external static magnetic field. For ^{43}Ca , the nuclear spin is $I = 7/2$ and the nuclear g-factor has been measured as $g_I = (2/7) \times -1.31537(60)$ [Ols72] in the neutral atom. There will be a small difference in the ion as there will be less diamagnetic shielding from the electrons, however this has yet to be measured experimentally. For the $S_{1/2}$ ground state the hyperfine splitting energy is $E_{hfs}/h = -3225.60828640(32)$ MHz [ABG⁺94] and the g-factor for the valence electron $g_J = 2.00225664(9)$ [TPR⁺03]. The eigenvectors of this solution give the α and β coefficients which we tabulate for magnetic fields of interest in appendix B.1.

6.2.2 Transitions

The magnetic dipole operator is $-\boldsymbol{\mu} \cdot \mathbf{B}_{\mu w}$ where $\mathbf{B}_{\mu w}$ refers to the oscillating magnetic field, so the transition matrix elements are given by

$$R_{mn} = -\langle m | \boldsymbol{\mu} \cdot \mathbf{B}_{\mu w} | n \rangle. \tag{6.3}$$

and the transition is driven with the Rabi frequency $\Omega_R = R_{mn}/\hbar$.

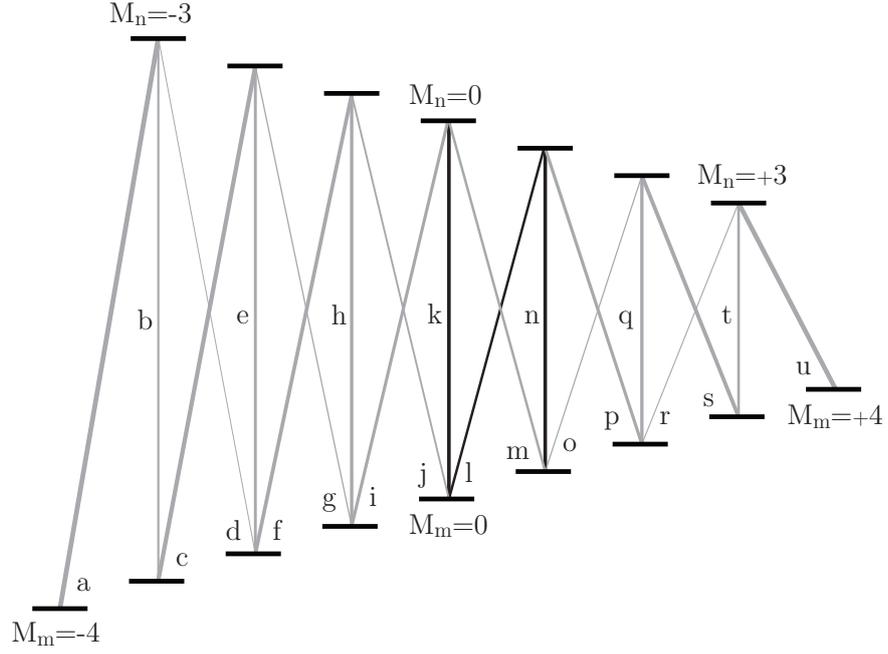


Figure 6.1: The hyperfine levels of the $S_{1/2}$ ground state of $^{43}\text{Ca}^+$. The line thicknesses are proportional to the sizes of the transition matrix elements. The three potential qubit transitions are coloured black.

We define the three components of $\mathbf{B}_{\mu\nu}$ as π , which drives $M_n = M_m$ transitions, and σ^\pm , which drive $M_n = M_m \pm 1$ transitions [Sta09]. This polarization is given by $q = M_n - M_m$. For a given q

$$R_{mn} = -\langle m | \boldsymbol{\mu} \cdot \mathbf{B}_{\mu\nu} | n \rangle = -(-1)^q B_q \langle m | \mu_{-q} | n \rangle \quad (6.4)$$

$$\begin{aligned} &= (-1)^{q+1} B_q \{ \alpha_m \alpha_n \langle -\frac{1}{2}, M_m + \frac{1}{2} | \mu_{-q} | -\frac{1}{2}, M_n + \frac{1}{2} \rangle \\ &\quad + \beta_m \beta_n \langle \frac{1}{2}, M_m - \frac{1}{2} | \mu_{-q} | \frac{1}{2}, M_n - \frac{1}{2} \rangle \\ &\quad + \alpha_m \beta_n \langle -\frac{1}{2}, M_m + \frac{1}{2} | \mu_{-q} | \frac{1}{2}, M_n - \frac{1}{2} \rangle \\ &\quad + \beta_m \alpha_n \langle \frac{1}{2}, M_m - \frac{1}{2} | \mu_{-q} | -\frac{1}{2}, M_n + \frac{1}{2} \rangle \}. \end{aligned} \quad (6.5)$$

We will ignore the interaction of the magnetic field with the nuclear moment here as the electronic moment is three orders of magnitude larger. This means the operator does not connect basis states with different M_I and we can ignore these terms. For π ($q=0$)

$$R_{mn} = -B_\pi \{ \alpha_m \alpha_n \langle -\frac{1}{2} | \mu_{-q} | -\frac{1}{2} \rangle + \beta_m \beta_n \langle \frac{1}{2} | \mu_{-q} | \frac{1}{2} \rangle \}, \quad (6.6)$$

for σ^+ ($q=1$)

$$R_{mn} = B_{\sigma^+} \alpha_m \beta_n \langle -\frac{1}{2} | \mu_{-q} | \frac{1}{2} \rangle, \quad (6.7)$$

and for σ^+ ($q=-1$)

$$R_{mn} = B_{\sigma^-} \beta_m \alpha_n \langle \frac{1}{2} | \mu_{-q} | -\frac{1}{2} \rangle. \quad (6.8)$$

The quantum numbers refer only to M_J as we have eliminated M_I since $\langle M_I | M_I \rangle = 1$. We now evaluate $\langle M_J | \mu_{-q} | M'_J \rangle$ using

$$\langle M_J | \mu_{-q} | M'_J \rangle = (-1)^{J+M_J+1} \begin{pmatrix} J & 1 & J \\ -M_J & -q & M'_J \end{pmatrix} \langle J | \mu | J \rangle \quad (6.9)$$

where

$$\langle J | \mu | J \rangle = -g_J \mu_B [J(J+1)(2J+1)]^{1/2}. \quad (6.10)$$

In our case $J = 1/2$ so by evaluating the required 3-j symbols we get

$$\pi : R_{mn} = \frac{1}{2} \mu_B g_J B_\pi (\alpha_m \alpha_n - \beta_m \beta_n) \quad (6.11)$$

$$\sigma^+ : R_{mn} = -\frac{1}{\sqrt{2}} \mu_B g_J B_{\sigma^+} \alpha_m \beta_n \quad (6.12)$$

$$\sigma^- : R_{mn} = \frac{1}{\sqrt{2}} \mu_B g_J B_{\sigma^-} \alpha_n \beta_m \quad (6.13)$$

Appendix B.2 contains tabulated matrix elements for all the transitions at the three static fields we consider in this chapter.

6.2.3 Potential qubits

In this section we discuss the differing features of three potential qubits we could use for our entangling gate.

6.2.3.1 Low-field qubit

The zero field clock transition between $M_m = +0$ and $M_n = +0$ (k in fig 6.1), whilst robust to magnetic field noise, has two major practical issues that prevent its use as a qubit at exactly zero field. The first is that at very low field the population will be coherently trapped in a dark state of the $D_{3/2}$ level when

laser cooling is attempted [BB02] (though polarization modulation schemes exist which can reduce this problem). The second is that because there is no Zeeman splitting the two sets of M states are degenerate and there is no way of spectrally resolving the different microwave transitions between states. Both of these problems can be solved by applying a small magnetic field. In previous work by the group a field of 1.78 G was used which gave a first-order magnetic field dependency of 4.33 Hz/mG [LKH⁺07]. The measured coherence time was $T_2 = 1.2(2)$ s and $\gtrsim 45$ s with a spin echo. A field of this size would shift non-qubit transitions (i, j, l and m in fig 6.1) to frequencies similar to those of the motional sideband frequencies we will be using for our entangling gate (see section 6.6) which would lead to errors in the gate through off-resonant excitations and ac Zeeman shifts (as seen in work at Innsbruck [KBZ⁺09b]). If we stipulate that no carrier transition should be closer to the motional sideband transition than the qubit carrier itself we find we require a field of 23 G. At this field the closest transition to the qubit is ~ 8 MHz away compared to the motional sidebands 4 MHz from the qubit (typical for our trap). The first order magnetic field dependence is increased to 55.9 Hz/mG but this should still give us a decoherence time several orders of magnitude longer than the entangling gate time. 23 G appears to be a good starting point but it could easily be adjusted to achieve the best compromise between qubit coherence time and suppression of off-resonant transitions.

6.2.3.2 Field-independent transition at 146 G

By differentiating the Breit-Rabi formula we can calculate the field-dependence of each transition as a function of magnetic field and thus find other transitions which have a field-independent point. The first one that occurs at non-zero field is the $M_m = 0$ to $M_n = +1$ transition (l in fig 6.1) at 146.094 G and the second is the $M_m = +1$ to $M_n = 0$ transition (m), at a slightly higher field of

146.331 G. Both l & m have the advantage over working at low field that we can operate exactly at the field-independent point. The coherence time should therefore be even better than the 1.2 s measured at 1.78 G. The splitting between adjacent microwave transitions is ~ 50 MHz so much faster initialisation pulses will be possible than at 23 G and off-resonant excitations and ac Zeeman shift from non-qubit transitions during the gate should be negligible. These transitions have two disadvantages though, firstly they are σ transitions and we can only generate linear microwave polarizations (circular polarization would require a $B_{\mu w}$ component in y but this is zero by symmetry, see section 7.1.3). This means that for a linear polarization perpendicular to the static field only half the field is driving the qubit. Secondly, the matrix elements of the transitions are 30% weaker than the other two qubits under discussion. Of the two transitions we would prefer l in practice as it requires one fewer initialisation transition to reach from the $M = \pm 4$ states (see section 6.4.1.2).

6.2.3.3 Field-independent transition at 288 G

This is the next field-independent transition. It occurs at 287.783 G and is from $M_m = +1$ to $M_n = +1$ (n in fig 6.1). Like the low-field qubit this is again a strong π transition and the Zeeman splitting between transitions is even higher (~ 100 MHz). The analogous transition in $^{25}\text{Mg}^+$ (at 213 G) is the transition that was used in the microwave gate demonstration [OWC⁺11].

6.2.3.4 Summary

There are three potential qubits of interest for our gate which are summarised in table 6.2. The higher field qubits offer definite advantages but at greater experimental effort due to possible difficulties in Doppler cooling them efficiently (see section 6.3 for discussion) and the requirement for large magnetic fields. It should become more clear which one is favoured as work towards a gate pro-

gresses.

Table 6.2: A summary of the transition properties of the three qubits discussed. f is the qubit frequency, and pol. gives the polarization of the microwave field required to drive the transition (where π indicates $\mathbf{B}_{\mu w} \parallel \mathbf{B}_0$).

B_0 (G)	f (GHz)	$ \downarrow\rangle$ M_m	$ \uparrow\rangle$ M_n	$\frac{df}{dB_0}$ (MHz/G)	$\frac{d^2f}{dB_0^2}$ (kHz/G ²)	Matrix element ($\mu_B B_{\mu w}$)	Pol.
23	3.226	0	0	55.9	2.43	1.00	π
146	3.200	0	+1	0	2.41	0.70	σ^+
288	3.123	+1	+1	0	2.51	1.00	π

6.2.4 AC Zeeman shift

In the case of an off-resonant microwave field, where the detuning $\Delta \gg \Omega_R$, the transition rate between states is negligible. There is a shift in the energies of the upper and lower levels (see figure 6.2) of

$$\pm \frac{\Omega_R^2}{4\Delta} \quad (6.14)$$

which follows from the two-state Bloch equations. For E1 transitions this shift is often called the ac Stark shift of ‘light shift’. When driving the gate there will be a shift of this type involving the carrier transition if there is a residual microwave field at the ion during the gate (see section 6.6). In this case $\Delta \approx \pm 2\pi \times 4$ MHz for the two sidebands. The shift from the red and blue sidebands will be equal and opposite in sign but they will not necessarily cancel as the two drive systems are separate (see section 7.2) and so the nulling error is not common-mode.

There will also be ac Zeeman shifts due to the four off-resonant transitions to states other than the qubit states. For the 146 G and 288 G qubits, these transitions are of order 50 and 100 MHz detuned respectively and so their effect will be small in comparison. For the 23 G qubit however we calculate the net shift from all four off-resonant transitions to be 67% (64%) of the shift due to the carrier for the red sideband (blue sideband) and opposite in sign. Note that this shift is caused by the component of $\mathbf{B}_{\mu w}$ orthogonal to the static field, whereas

the shift on the carrier is from the component of $\mathbf{B}_{\mu w}$ parallel to the static field so these errors contribute independently.

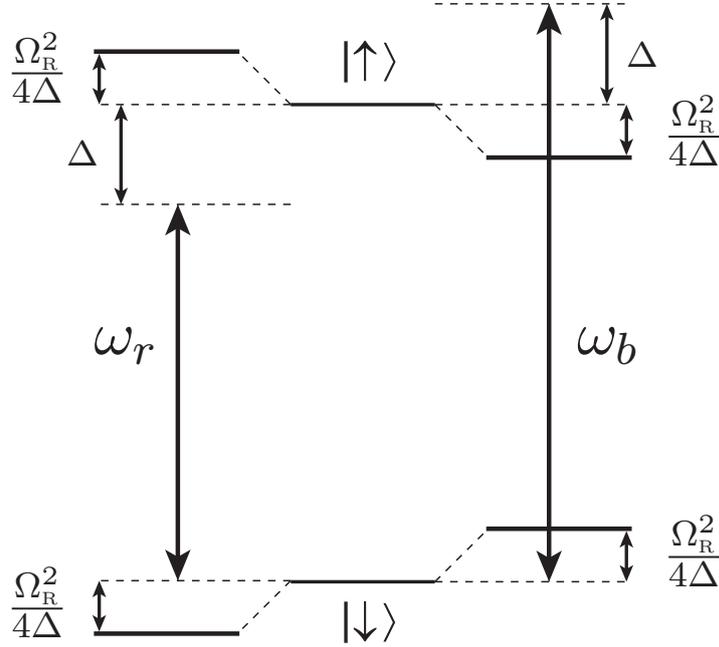


Figure 6.2: The effect of off-resonant microwave fields on the qubit states' energies. The left side shows the effect of a red-detuned field and the right shows the effect of a blue-detuned field.

6.3 Doppler cooling

Laser cooling is required to initialise the ion crystal's motional state. The entangling gates we are considering are broadly insensitive to the ions' initial motional state (see section 6.6) and it will suffice to reach a thermal state close to the Doppler limit ($\bar{n} = 2.8$ for a 4 MHz motional mode). This eliminates the need for more complicated sub-Doppler cooling schemes and removes a large experimental overhead.

$^{43}\text{Ca}^+$ has already been successfully Doppler cooled by our group at fields of a few Gauss using a 397 nm laser with 3.2 GHz sidebands and a single 866 nm repumping frequency. Fluorescence rates are typically only $\sim 25\%$ of that achieved in $^{40}\text{Ca}^+$ due to coherent population trapping in dark states [BB02]. At

23 G the Zeeman splittings are still of order the natural linewidths of the laser transition used and so we expect qualitatively the same cooling scheme to work. Due to the larger field we expect a quicker procession out of any dark states and thus a possible improvement in the fluorescence rate.

No group has yet attempted to laser cool an odd isotope ion with low-lying D states at intermediate field. The two intermediate-field qubits hitherto demonstrated, $^9\text{Be}^+$ [LOJ⁺05] and $^{25}\text{Mg}^+$ [OWC⁺11], have the advantage that they have a relatively simple level structure that has closed laser-cooling transitions (from $S_{1/2} M = I \pm \frac{1}{2}$ to $P_{3/2} M = I \pm \frac{3}{2}$). In $^{43}\text{Ca}^+$, decays to the D level will cause population to spread out to different $S_{1/2}$ states. At low field this is not so much of a problem as none of these states is more than a few MHz detuned from the cooling laser and will be efficiently repumped. At intermediate field though these states can be hundreds of MHz away from the cooling laser frequency which would cause population to accumulate there, greatly reducing the fluorescence rate. A possible scheme to deal with this problem is outlined below. Despite the difficulties there are two compelling advantages in using $^{43}\text{Ca}^+$. The first is that the low-lying D-states allow for electron shelving which permits high readout fidelity [MSW⁺08]. The second is that only the ions with low-lying D-states (Ca^+ , Sr^+ and Ba^+) have visible laser cooling transitions. Yb^+ is an intermediate case in that it has a relatively long wavelength and low-lying metastable states but electron shelving has not been demonstrated due to an unfavourable level structure. Visible cooling transitions are important as fibre optics that can be integrated into the trap and high finesse cavities perform poorly at shorter wavelengths. Visible light is also less likely to charge the trap [WLL⁺11] and cause drifts in the ions' secular frequencies, a major source of error in work with Mg^+ [OWC⁺11].

6.3.1 S-P-D level structure

Figure 6.3 shows the transitions in $^{43}\text{Ca}^+$ at low magnetic fields and their isotope shifts compared to $^{40}\text{Ca}^+$. Figures 6.4 to 6.8 show the hyperfine and Zeeman splittings in all the low-lying levels of $^{43}\text{Ca}^+$ at fields of up to 300 G.

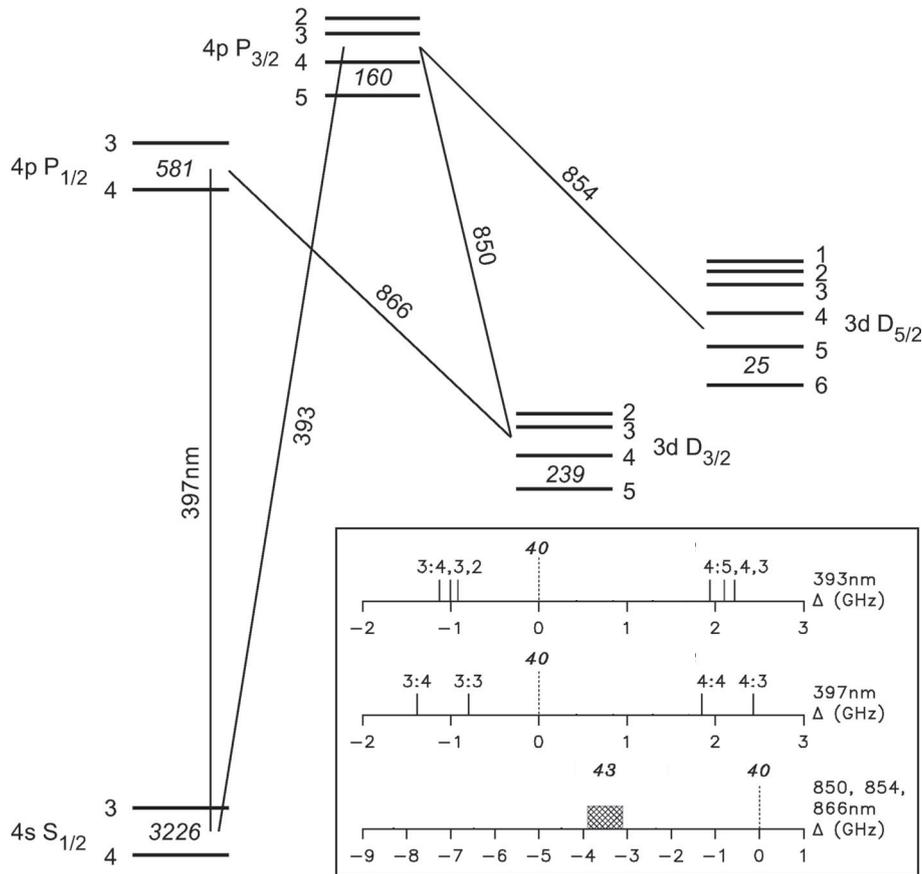


Figure 6.3: The transitions in $^{43}\text{Ca}^+$ at low magnetic fields. The largest hyperfine splitting in each term is indicated in MHz. Inset: Positions of the hyperfine components of the 393 nm and 397 nm transitions (labeled with $F_{lower} : F_{upper}$) compared to $^{40}\text{Ca}^+$. The isotope shifts in the infrared transitions (similar in all three transitions) are also shown, without the hyperfine details. Figure modified from [LRH⁺04].

6.3.1.1 Intermediate field cooling scheme

The group is currently conducting a numerical study of $^{43}\text{Ca}^+$ using the optical Bloch equations [SLA⁺08] to ascertain what the best laser parameters are for cooling in the intermediate field. Initial results indicate that high fluorescence is

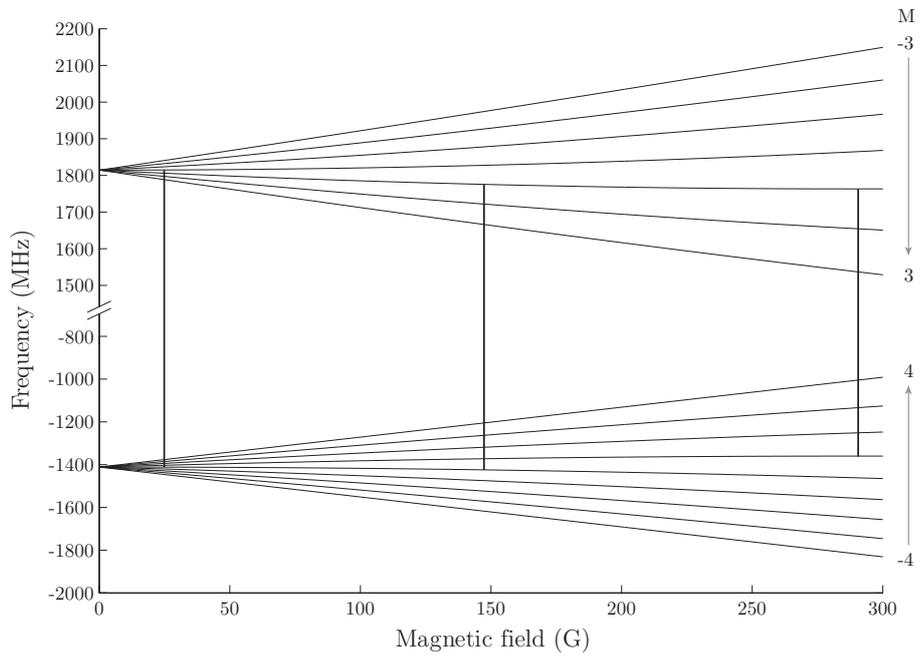


Figure 6.4: The hyperfine states of the $S_{1/2}$ ground level of $^{43}\text{Ca}^+$ as a function of magnetic field. The vertical lines mark the three qubit transitions discussed.

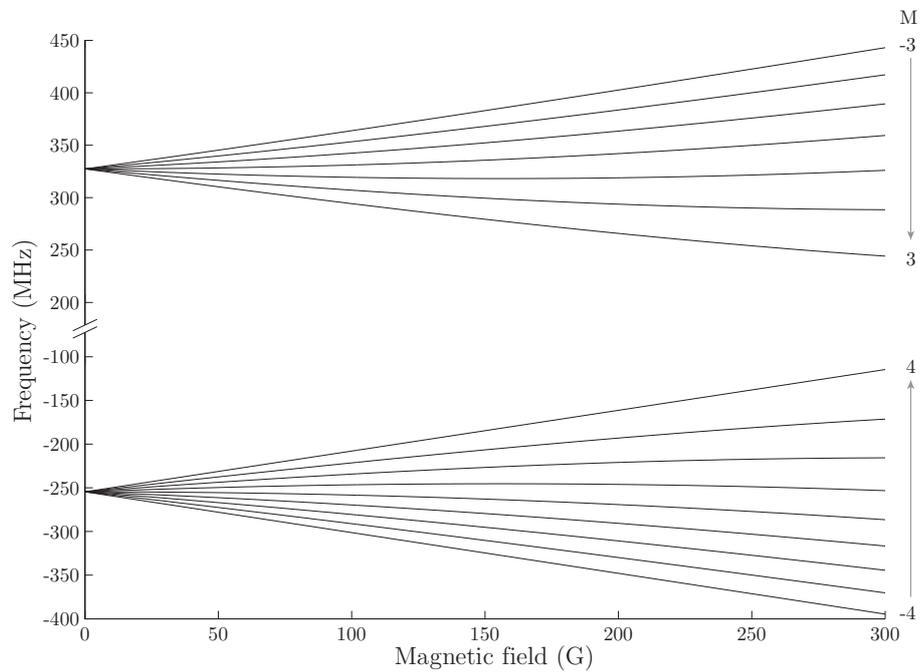


Figure 6.5: The hyperfine states of the $P_{1/2}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field.

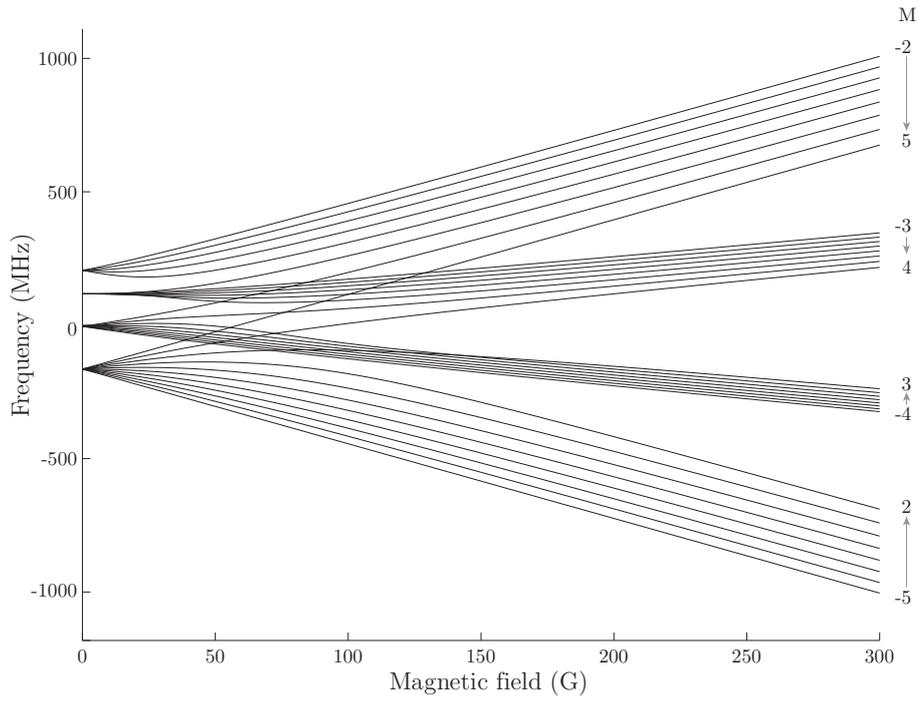


Figure 6.6: The hyperfine states of the $P_{3/2}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field.

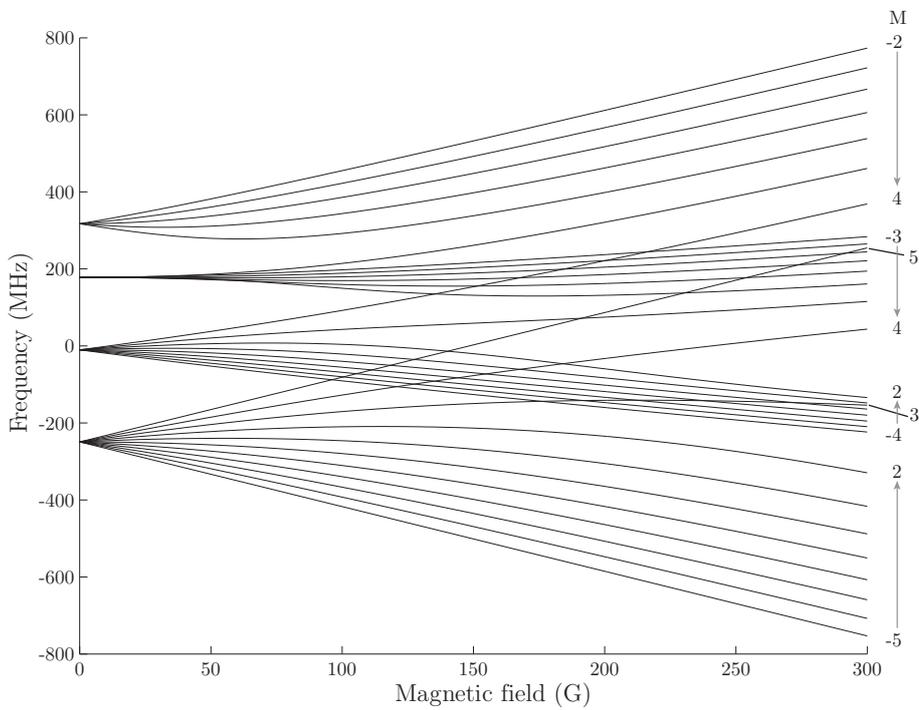


Figure 6.7: The hyperfine states of the $D_{3/2}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field.

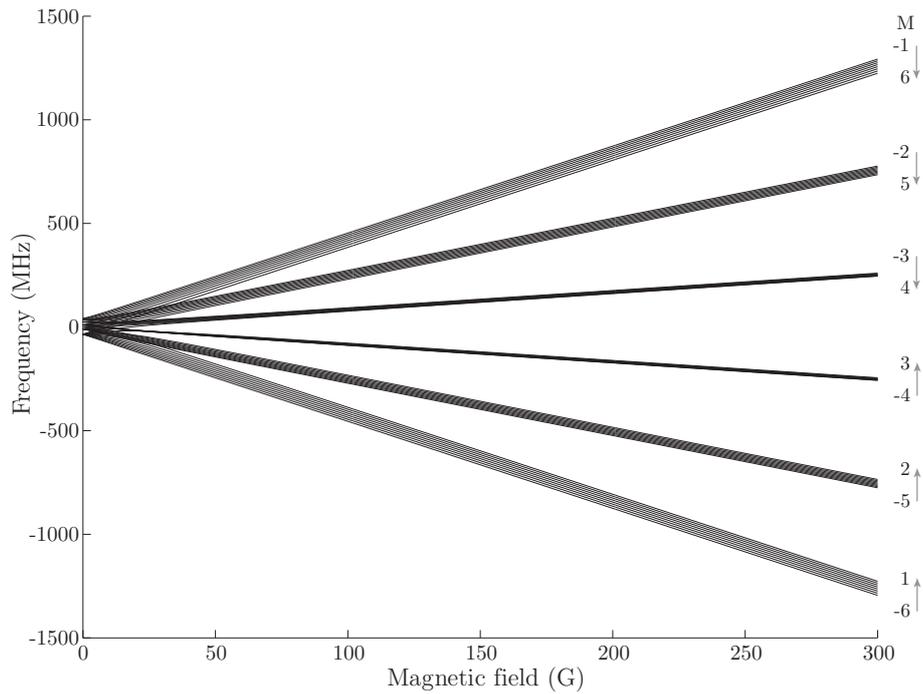


Figure 6.8: The hyperfine states of the $D_{5/2}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field.

achievable by using only two 397 nm frequencies (as at low field). This is possible if 397 nm laser beams of predominantly only one σ polarisation of light are used to ensure only a few states at one end of the hyperfine manifold are populated. Figure 6.9 shows a simulated scan of the 397 nm laser at 146 G. The peak population of 0.16 compares favourably with the peak population of 0.25 achievable in $^{40}\text{Ca}^+$ and the linewidth is narrow enough to permit cooling to around 3 times the Doppler limit. Further optimisation and experimental verification are currently underway.

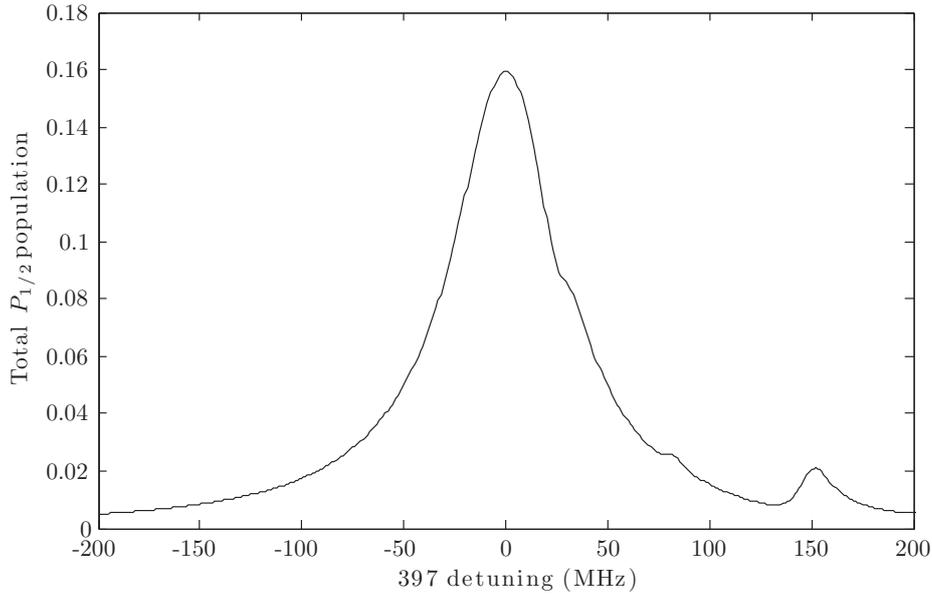


Figure 6.9: Simulated total $P_{1/2}$ population as a function of the 397 nm laser detuning (scale offset so the peak is at zero detuning) at 146 G. There are two 397 nm lasers each with $I = 20I_s$ and offset by 2.931 GHz from each other. The 397 nm beam's ratio of $\sigma^+ : \pi : \sigma^-$ polarisations is 0.87:0.13:0. There is a single 866 nm laser with $I = 1000I_s$. The 866 nm beam's ratio of $\sigma^+ : \pi : \sigma^-$ polarisations is 0.03:0.53:0.44.

6.4 Single-qubit operations

6.4.1 Qubit initialisation

6.4.1.1 Optical pumping to $M = \pm 4$

We can easily prepare the ion in either the $M = +4$ or $M = -4$ states by applying σ^+ or σ^- polarised 397 nm light only (as well as the 866 nm repumper). This is because this laser couples to all states except this one so any population which decays into it from the $P_{1/2}$ is trapped there. The fidelity of this preparation is limited only by the polarization impurity in the 397 nm beam (assuming one waits long enough to reach the steady state). We will pump to whichever of the $M = \pm 4$ states is already significantly populated by the Doppler cooling process (see above) to ensure the steady state is achieved rapidly.

6.4.1.2 Microwave transfer pulses

The qubit is initialised by transferring the population optically pumped to one of the $M = \pm 4$ stretch states to one of the qubit states using a series of microwave transitions. Initially we will use simple π pulses. The population of the destination state after a microwave pulse is given by

$$\frac{\Omega_R^2}{\Omega^2} \sin^2 \left(\frac{\Omega t_{pulse}}{2} \right) \quad (6.15)$$

where Ω_R is the Rabi frequency, t_{pulse} is the pulse time and $\Omega = \sqrt{\Omega_R^2 + \Delta^2}$ where δ is the frequency detuning of the pulse from resonance. For a π -pulse, $\Delta = 0$ and $t_{pulse} = \pi/\Omega_R$. This assumes that the Rabi frequency is small enough that off-resonant excitations on other transitions can be ignored.

As the transitions used for these initialisation pulses are field-dependent these π -pulses are not robust to magnetic field drifts; however initialisation fidelities of over 99% should be possible with our initial setup. For $\Omega_R = 2\pi \times 100$ kHz this will require a microwave power stability of $\pm 6\%$, which is readily available from a synthesiser, and a static field stability of ~ 10 mG (measured lab drifts over a day are of this size [Kei07]). Apart from improving microwave power stability and reducing magnetic field noise and drift there are two other methods of improving the initialisation fidelity. Firstly we could use adiabatic pulses [Sch04]. Adiabatic pulses would require extra pulse-shaping and frequency chirping electronics though and the initialisation time would be longer. Secondly we could use composite pulses [Lev86] where the phase changes in discrete steps during the transfer pulse. This allows transfers which are more robust than a π -pulse to off-resonant excitation, amplitude fluctuations and frequency drift. It would however require a frequency source whose phase can be rapidly adjusted.

6.4.2 Single-qubit gates

Any quantum algorithm can be broken down into a sequence of operations that only uses one specific two-qubit entangling operation plus a number of single-qubit rotations [NC00]. We will use microwave carrier pulses to drive single qubit rotations of the form

$$R(\theta, \phi) = I \cos(\theta/2) + i(\sigma_x \cos \phi - \sigma_y \sin \phi) \sin(\theta/2). \quad (6.16)$$

where I is the identity matrix and σ_x and σ_y are the Pauli-spin matrices. θ is set by the length of the pulse and ϕ by the phase of the pulse. Section 7.2 discusses how these pulses are generated. We will not be able to address different ions individually with microwave pulses although future experiments should add this capability (see chapter 9).

Microwave driven single qubit gates have recently been performed with fidelities exceeding the fault-tolerant threshold. The NIST group reports errors of $2.0(2) \times 10^{-5}$ [BWC⁺11] in the ${}^9\text{Be}^+$ clock qubit at 15.1 G (95 Hz/mG field sensitivity).

6.5 Qubit readout

6.5.0.1 Microwave transfer pulses

The population of one of the two qubit states is transferred back to either of the $M = \pm 4$ states by running the initialisation pulse sequence (see section 6.4.1.2) in reverse.

6.5.0.2 Qubit shelving

We shelve the qubit population now in $M = \pm 4$ by applying 393 nm σ^\pm polarised light to excite the transition to the $P_{3/2}$ $M = \pm 5$ state. 6.5% of the time this state will decay to a D-state but if not it can only decay back to the $S_{1/2}$

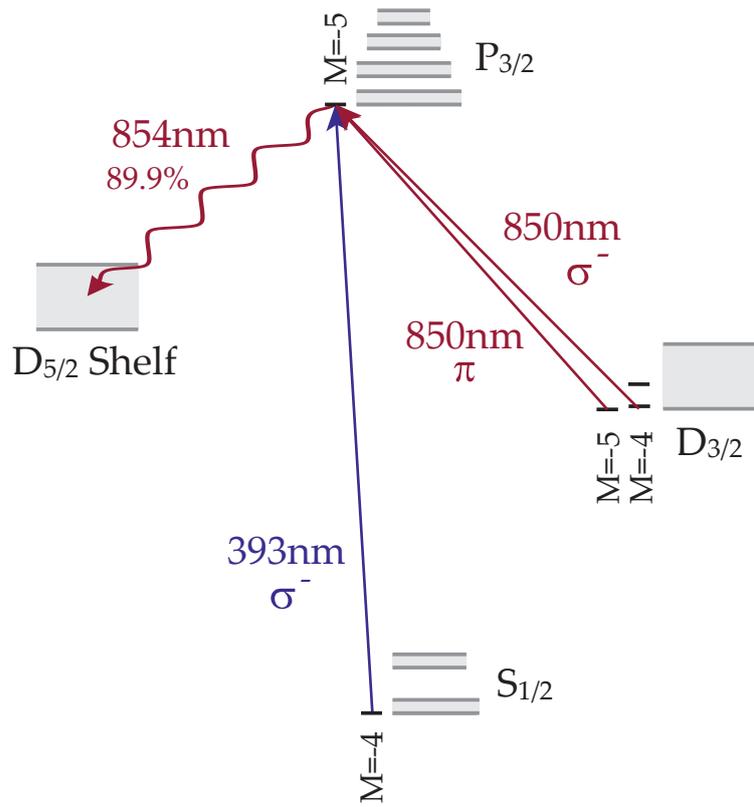


Figure 6.10: Shelving scheme for qubit readout. The the same scheme can be applied with σ^+ 393 nm and 850 nm to shelve the $S_{1/2}$ $M = +4$ state.

$M = \pm 4$ where it can be excited again. 89.9% of the D-state decays will be to a state in the metastable $D_{5/2}$ level, at which point the ion is shelved. The 10.1% that decays to the $D_{3/2}$ level can then be pumped back to the $P_{3/2}$ $M = \pm 5$ using the 850 nm laser and another attempt made to shelve using the 393 nm. The 850 nm laser polarization is chosen such that the intensity of the σ^\pm component is much higher than the π component to ensure the $M = \pm 4$ states are preferentially pumped to the $M = \pm 5$ state. The sources of error for shelving are off-resonant excitation of the other qubit state by the 393 nm light, incorrect polarizations of the 393 nm and 850 nm lasers that cause leakage to other states and decay from $D_{5/2}$. A thorough numerical optimisation of the best laser parameters [Szw09] yielded a fidelity for shelving of 99.968% in 100 μs . Experimentally 99.77(4)% was achieved in 400 μs . For the higher field qubits we expect the re-

sults to be broadly similar but with the possibility for fidelity gains due to the frequency discrimination that the large Zeeman splittings provide.

One could also shelve the qubit directly to the $D_{5/2}$ level using the electric quadrupole transition at 729 nm [BKRB08a]. Because this transition is very narrow (~ 1 Hz) off-resonant excitation of the other qubit state is avoided. Also, if this pulse has imperfections it can be repeated many times by shelving to different states of the $D_{5/2}$ level to improve the fidelity. The error associated with $D_{5/2}$ decay is still present though and so the overall fidelity improvement is marginal. The narrow laser required for addressing the quadrupole transition is technically demanding, whereas the 393+850 nm solution is experimentally simpler and more scaleable.

6.5.0.3 Photon counting

Once one qubit state is shelved to the $D_{5/2}$ level, measurement is carried out by turning on the 397 nm and 866 nm Doppler cooling lasers and detecting the fluorescence photons. As the $D_{5/2}$ state does not fluoresce the qubit state can be determined with very high fidelity. Our group has performed a detailed study of how well this can be done, both with a single ion using a PMT and multiple ions using an EMCCD [Bur10]. For characterising the fidelity of a single two-qubit gate we do not require single-shot readout and we can use many experiments to build up a photon number histogram from the PMT [SKK⁺00].

6.6 Multi-qubit entangling operations

6.6.1 Driving motional sidebands with microwaves

Laser fields can be used to simultaneously drive transitions between both the internal states of a trapped ion and its external motional states [WMI⁺98a]. These sideband transitions can entangle the internal and external states and

are the mechanism by which multiple ions can be entangled via their shared motional states. For cold, tightly-bound ions (the Lamb-Dicke approximation) the Rabi frequencies for these sidebands (see figure 6.11) are $\Omega_r = \sqrt{n}\eta\Omega$ and $\Omega_b = \sqrt{n+1}\eta\Omega$ where Ω is the Rabi frequency of the carrier. The Lamb-Dicke parameter is given by $\eta = k_j q_0^j$ where k_j is the projection of the laser field's wavevector in the direction of the motional mode j and $q_0^j = \sqrt{\hbar/(2m\omega_j)}$ is the spread of the zero-point wave function. For the 729 nm quadrupole transition in $^{43}\text{Ca}^+$ and $\omega_j = 2\pi \times 4\text{MHz}$, $\eta \approx 0.05$. For the 3.2 GHz microwave transition however, $\eta \approx 4 \times 10^{-7}$, which makes the sideband frequencies unfeasibly slow for gate operation.

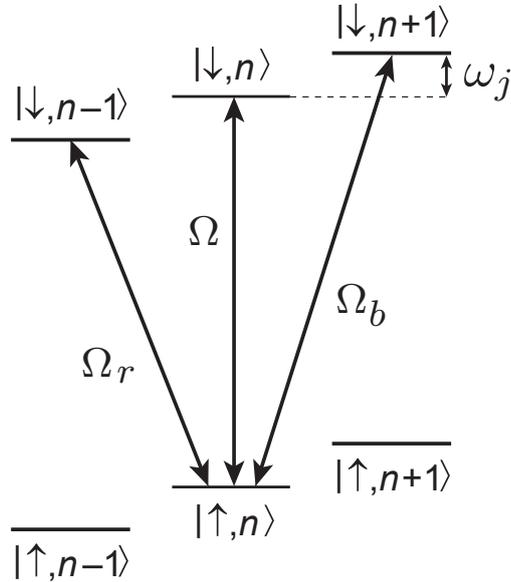


Figure 6.11: A diagram of the levels, driving fields and the relevant detunings involved in driving the carrier and motional sideband transitions of a single ion.

This poor coupling to the motion is because the microwave field varies so little across the ion's wavefunction compared to in an optical field. However, Ospelkaus et al. [OLA⁺08] have shown that in the near-field of a microwave current carrying conductor the gradients are much stronger than in a free-space wave and sidebands can be driven at frequencies comparable those achieved with a laser field. Unlike a laser field, this near-field does not necessarily have

a simple analytical form so the expression for the Lamb-Dicke parameter no longer holds. We can however calculate an ‘effective Lamb-Dicke’ parameter η' by taking the ratio of the carrier ($\propto B_{\mu w}$) to the sideband Rabi frequency ($\propto B'_{\mu w} q_0^j$) for an ion above a single current carrying wire (where $B_{\mu w} \propto 1/d$). For $^{43}\text{Ca}^+$ trapped $d = 75 \mu\text{m}$ above this wire this number is $\eta' \approx 4 \times 10^{-5}$. Whilst this is large enough to allow a gate in a reasonable time, unwanted off resonant excitations and ac Zeeman shifts from the strong carrier would dominate. By using a more complex arrangement of wires however the absolute value of the B-field at the ion can be made very small but the strong gradient retained. The required trap geometry is discussed further in section 7.1.3.

6.6.1.1 Field gradient directions

Figure 6.12 shows the magnetic field vectors in the region around the ion in the case where $B_{\mu w}$ has been nulled at the ion. For obtaining the maximum gate speed it is important that the motional mode used for driving the gate is aligned with the direction of greatest microwave field gradient. We choose the static B-field, \mathbf{B}_0 to be along x : it must be parallel to the substrate in order to allow σ -polarised laser beam access for state preparation and readout (see section 8.7).

The 23 G and 288 G qubits are π transitions are driven by the x -component of the microwave field, $B_{\mu w}^x$. The motional sidebands of the qubit transition are driven by the gradient of $B_{\mu w}^x$, which lies in the z direction, $B'_{\mu w} = \partial B_{\mu w}^x / \partial z$. Thus the motional mode should be aligned with z for maximum gate speed. Conversely, the 146 G qubit is a σ transition, driven by $B_{\mu w}^z$ so the motional mode should lie aligned with x . In practice however the modes will have to be rotated by a small angle from this otherwise the mode along z would not be efficiently Doppler-cooled because it is perpendicular to the cooling beam direction (1,1,0). This rotation can be achieved by applying a differential dc voltage to the two rf electrodes and the dc electrodes as discussed in section 3.1.

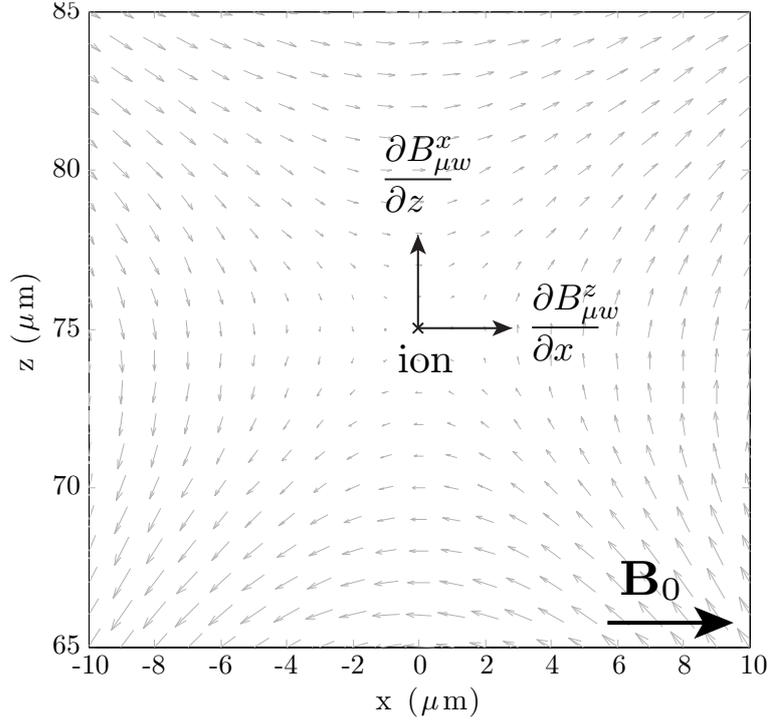


Figure 6.12: The magnetic field vectors in a $20 \times 20 \mu\text{m}$ square around the ion. The surface of the trap electrodes is at $z = 0$. The direction of the static field \mathbf{B}_0 and the gradients of the x and y components of the microwave field $\mathbf{B}_{\mu w}$ at the ion are indicated.

6.6.2 The Mølmer-Sørensen gate

In order to entangle our ions we will use a Mølmer-Sørensen gate [MS99, SM99, SM00]. This gate works by inducing correlated spin flips in a string of n ions. The Hamiltonian of the gate for the two ion case is

$$H \propto \sigma_\phi \otimes \sigma_\phi \quad \text{where} \quad \sigma_\phi = \sigma_x \cos \phi + \sigma_y \sin \phi. \quad (6.17)$$

The Hamiltonian is realised by exciting both ions simultaneously with a bichromatic microwave field with frequencies detuned by δ from the red and blue ion crystal vibrational mode sidebands ($\omega_{b,r} = \omega_0 \pm \omega_j \mp \delta$). ϕ is defined by the phase of the two sidebands as $\phi = (\phi_b + \phi_r)/2$. Figure 6.13 shows diagrammatically how these collective spin flips are achieved.

After a time $\tau = 2\pi/\delta$ the motional mode will be returned to its initial state

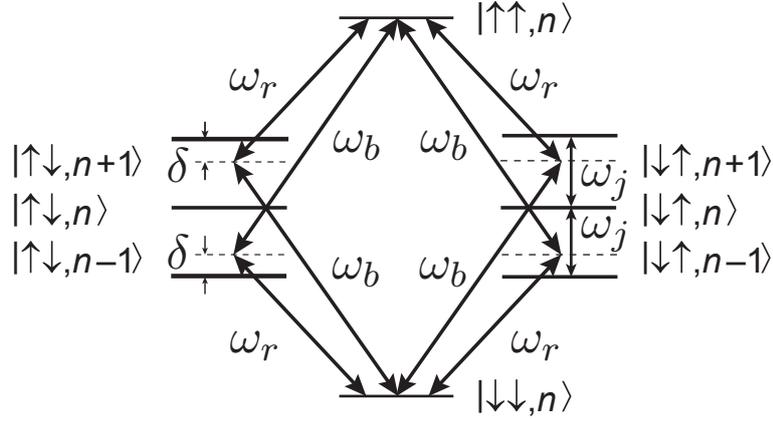


Figure 6.13: A diagram of the levels, driving fields and the relevant detunings involved in the Mølmer-Sørensen gate. Modified from [Roo08].

and a gate will have been carried out with the propagator

$$\mathcal{U}_{\varphi\varphi}(\tau) = \exp\left[\frac{2\pi i}{\delta^2} \left(\sum_n \Omega_{j,n} \sigma_{\varphi_s}^n\right)^2\right] \quad (6.18)$$

where $\Omega_{j,n}$ is the sideband Rabi frequency for ion n on the motional mode j .

This Rabi frequency is given by

$$\Omega_{j,n} = \frac{b_{j,n} q_0^j B'_{\mu w}}{2\hbar} \mu_{\uparrow\downarrow} \quad (6.19)$$

where $b_{j,n}$ is the normal-mode coefficient [Jam98], $q_0^j = \sqrt{\hbar/(2m\omega_j)}$ is the spread of the zero-point wave function and $\mu_{\uparrow\downarrow}$ is the transition matrix element. $B'_{\mu w}$ is the gradient, in the direction of the mode j , of the $\mathbf{B}_{\mu w}$ component that drives the transition we are considering (see previous section).

The advantage of this gate scheme is that the gate's interaction strength is very insensitive to the ions' motional state n and so n can be a thermal state.

6.6.3 Gate speed

The gate is driven in $t_{gate} \approx \frac{\pi}{2} \Omega_{j,n}$. For a two-ion $^{43}\text{Ca}^+$ gate driven on a 4 MHz transverse rocking mode ($b_{j,n} = [-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}]$) the gate time is given by

$$t_{gate} \approx \frac{1}{\mu_{\uparrow\downarrow} B'_{\mu w}} \times 9.32 \text{ ms} \quad (6.20)$$

where $\mu_{\uparrow\downarrow}$ is in units of μ_B . Therefore using either the 23 G or 288 G qubit a field gradient of 9.32 T/m will give us a gate time of 1 ms and we require $\delta = 2\pi \times 1$ kHz. The rocking mode is chosen instead of the COM mode as only the gradient of the electric field noise causes heating on this mode, greatly suppressing the motional heating of the ions compared to the COM mode [KWM⁺99].

6.6.3.1 Advantages over laser-driven gates

Driving the gate with microwaves has several advantages over the laser-driven gates that have been widely used over the last decade.

For gates driven by Raman laser transitions large detunings have to be used to reduce off-resonant photon scattering. This requires hundreds of mW of power and in Ca^+ there is a fundamental limit on the minimum error achievable of 1.06×10^{-4} due to scattering to the D-states [OIB⁺07, UBV⁺10]. For gates driven by quadrupole transitions photon scattering is not an issue, but a very high quality laser (< 1 Hz linewidth) is required [BKRB08b]. In both cases a major limiting factor in achieving high fidelity is the technical difficulty of controlling the intensity, frequency and phase of the laser field at the ions. The required stabilities in microwave systems are typically available off the shelf and if not then a large body of literature and expertise exists such that the required systems can be easily developed. The laser systems used on the other hand are typically one-off systems that already represent the state of the art.

Whilst the Mølmer-Sørensen gate has been demonstrated on thermal states ($\bar{n} = 20(2)$) at the Doppler limit with lasers [KBZ⁺09a] there is a nonlinearity in the driving laser field because the ions' motional wavefunction is typically only about an order of magnitude smaller than the wavelength of the light. This is less of an issue in the microwave gates as the field gradient is linear to a very good approximation over several microns.

For microwave-driven gates the size of the trap is no longer limited by the diffraction limit of the gate laser beams. As you get closer to the surface the trap the microwave field increases as $1/d$ but the field gradient increases as $1/d^2$. This means as the trap gets smaller the gate speeds for a given current density and microwave power increase rapidly and nulling the microwave field at the ion gets relatively easier. Further miniaturisation will require a solution to the anomalous heating issue though (see chapter 5).

6.6.3.2 Experimental issues

Despite these advantages, there are several disadvantages that need to be overcome.

The ion must be placed at the null of the oscillating magnetic field with high accuracy. If the ion or the field drifts the resulting field will cause off-resonant carrier excitation and ac Zeeman shifts.

Generating the high magnetic field gradients is technically demanding due to the large microwave currents required and the heat this will generate. This means that initially, until traps can be miniaturised further, the gate will be relatively slow and therefore the detuning from the motional sideband δ is small. This makes the gate more susceptible to drifts in the motional mode frequency ω_j . Reversing the phase of one of the drive fields half way through the gate is an effective method of suppressing decoherence caused by this drift [HCD⁺11].

For a large array of ion traps there will be crosstalk between different gate regions. This is discussed further in chapter 9.

6.6.4 Alternative schemes

The use of microwave field gradients to drive entangling operations was also proposed independently by Chiavarini and Lybarger [CL08]. In this scheme a static 2-dimensional array of singly trapped ions was proposed with a mi-

crowave driving coil under each ion. The resultant trap would be ideally suited to direct simulation of coupled-spin systems. Whilst we will not discuss this scheme further some of the microwave engineering methods discussed in the rest of this thesis could equally be applied to realising this scheme.

There is also a long-standing proposal to use free-space microwaves to drive entanglement [MW01]. In this scheme field-dependent states are used and a strong static field gradient is applied to the ion crystal. The static gradient thus provides the state-dependent momentum transfer rather than the microwaves. The gradient also provides qubit frequency addressing due to each qubit being at a different frequency. This scheme has a few drawbacks though. Firstly the use of field-dependent states reduces the qubit coherence time. By using microwave-field dressed states the coherence time can be extended but at the cost of experimental overhead and with coherence times still well below that of a well-chosen clock state [TBJ⁺11]. The second issue is that as each qubit has a different, and position-dependent frequency, tracking of all the qubit phases during a computation may be complicated, especially if ions are to be shuttled around the trap. Finally the high gradients are technically challenging to produce and will only allow for 1-dimensional ion addressing. This still appears to be a promising route though as, if these issues can be overcome, the ion traps required will be much simpler than in the scheme we are pursuing.

7

Design of trap and electronics for microwave-driven quantum logic

This chapter assumes a working knowledge of basic rf and microwave engineering. Unless another source is explicitly referenced then a description of any component or technique used can be found in Pozar's standard text on microwave engineering [Poz05] or similar.

7.1 Trap design

7.1.1 Design goals

The trap design should meet four main design goals:

- Linear Paul trap with 1 MHz axial and 4 MHz radial trapping frequencies.
- Incorporate microwave current-carrying structures of a geometry that can produce a significant microwave field gradient $B'_{\mu w} = \partial \mathbf{B}_{\mu w}^x / \partial z = \partial \mathbf{B}_{\mu w}^z / \partial x$ at the ion but with $|\mathbf{B}_{\mu w}| = 0$.

- A $B_{\mu w}$ gradient of 10 T/m (to give a 0.93 ms two-qubit gate time) should be achievable with microwave powers in the few Watts regime.
- Impedance match the trap to 50 Ω drive electronics.
- Thermally design the trap so that it does not heat up excessively when microwaves are applied.

7.1.2 Trapping electrodes

We will first consider the parameters of the rf trap we will need for the confinement of the ion. For this project we will use a surface electrode trap. Such a geometry will allow planar microwave circuitry to be easily combined with the trap electrodes and placed directly under the ion. Surface trap technology also allows a route to future scalability (see chapter 9). In chip traps for neutral atoms this close integration of a microwave conductor on a surface with the current-carrying wires necessary for trapping atoms has already been demonstrated [Tre08].

7.1.2.1 Ion height

The first parameter we fix is ion-to-trap distance, d . As far as the entangling operation is concerned we want to be as close to the trap as possible because the magnetic field gradient (and thus the gate speed) has a d^{-2} dependence [OLA⁺08]. As the gate gets faster more operations can be performed before the qubit decoheres. Also, the detuning from the sidebands gets larger; this is important as it makes the gate less sensitive to fluctuations in the frequencies of the sidebands which can be a major error. Nulling the microwave field at the ion becomes relatively easier as the microwave fields scale more weakly, as d^{-1} . There are two issues that prevent us using a very small d . Firstly, motional heating increases with a $\sim d^{-4}$ dependence (see chapter 5), which means that the amount of heating during the gate will scale as $\sim d^{-2}$. Secondly, the ion should

be far enough from the surface to enable laser access. From figure 2.7 the minimum distance is around $50\ \mu\text{m}$ ($30\ \mu\text{m}$ has been achieved in Mg^+ [OWC⁺11] but there the laser wavelength is 279 nm rather than 397 nm). In the future this requirement could be circumvented by having two trapping regions and shuttling between them as required. One region would have a small d for gates and the other would allow laser access by either having a larger d or integrating optics into the trap itself [BEM⁺11].

It is hard to predict exactly where the best compromise will be until the factors limiting fidelity are better understood. Also, due to the trap-to-trap variability, a given trap's heating rate is generally not known until it is measured. We decided initially to use $d = 75\ \mu\text{m}$, a factor of two reduction compared with the prototype surface trap in chapter 3.

7.1.2.2 Trap parameters

Once the ion height is fixed, the widths of the two rf electrodes relative to the rf ground electrode between them is constrained by an analytical expression [Hou08] which is plotted in figure 7.1 with the assumption that the trap is symmetrical about the centre. We also plot the required rf voltage and resulting trap depth at a radial secular frequency of 4 MHz and stability parameter of $q = 0.3$. A centre electrode width of around $50\ \mu\text{m}$ gives the most efficient trap however the final design will use $70\ \mu\text{m}$ as this was found to be preferable for the microwave design. The expressions used are exact in the limit of no gaps between electrodes but the error introduced by the gaps is only a few percent if they are accounted for correctly [Sch10]. Symmetrical traps have slightly less desirable trapping parameters than asymmetrical ones (see section 3.1) but the design and analysis of the trap will be greatly simplified.

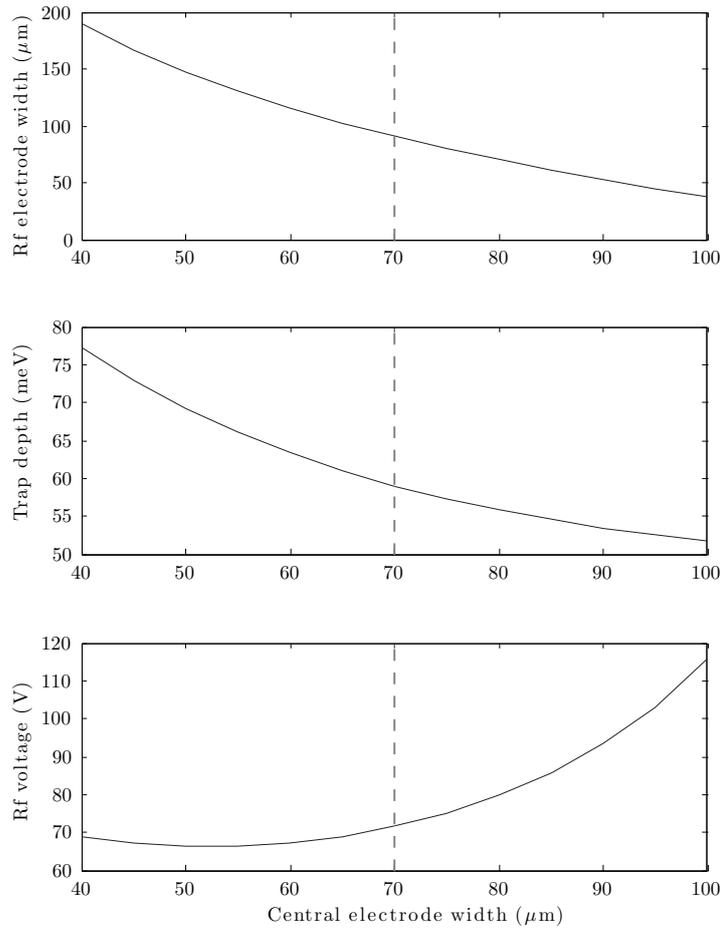


Figure 7.1: Trap parameters of a symmetric surface electrode trap as a function of electrode widths. The widths exclude gaps which are taken as $10 \mu\text{m}$. The ion is fixed at $d = 75 \mu\text{m}$ from the surface, the radial trapping frequency is fixed at $\omega = 2\pi \times 4 \text{ MHz}$ and the stability parameter is fixed at $q = 0.3$ (the choice of ω_r and q fixes the rf drive frequency at $\Omega_{r,f} = 2\pi \times 37.7 \text{ MHz}$). The $70 \mu\text{m}$ central electrode width of the final design is marked on the graphs.

7.1.3 Microwave conductors

The microwave conductors on the trap need to run parallel to the trap axis (see figure 7.2). This is so that when the field is nulled at one ion in a linear crystal of ions it is also nulled at all the other ions. In order to provide enough degrees of freedom to null $\mathbf{B}_{\mu w}$ in both the x and z (radial) directions at the ion we require three conductors and the ability to control the phase and amplitude of the microwave currents in each. The translational symmetry of the trap will ensure there is no component of $\mathbf{B}_{\mu w}$ along the axial direction y .

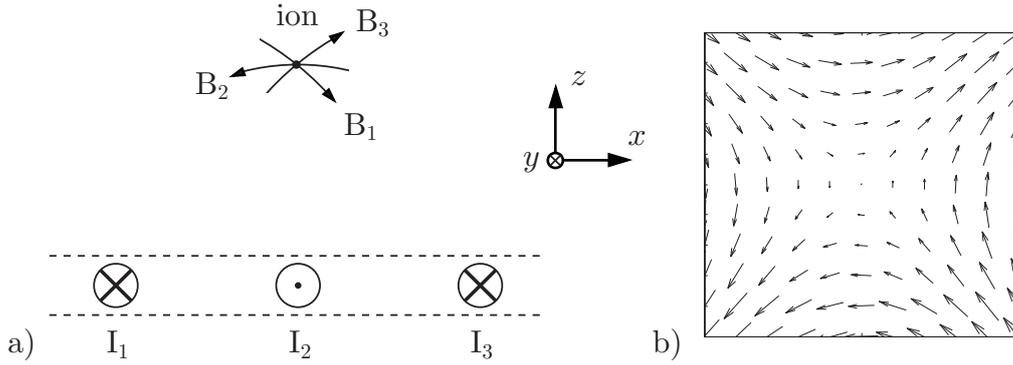


Figure 7.2: a) Layout of microwave currents on the trap required to null field at ions. b) Resultant $B_{\mu w}$ in the region around the ion.

7.1.3.1 Layout of conductors

It is desirable for the trap design to be compatible with our existing fabrication technique (see section 2.4) as developing a new fabrication technique would require significant time and effort. This limits us to having all the electrodes in a single plane, with no overlaps or vias. As we require two rf electrodes in a fixed position there are a limited number of choices for the layout of the three microwave conductors (see figure 7.3). Layout 2 was that originally proposed by Ospelkaus et al. [OLA⁺08] and layout 4 was the layout later experimentally realised [OWC⁺11]. In order to ascertain which layout is the most favourable we carried out a finite element simulation of each one (see 7.1.5.2). Finite element simulations are required because, at microwave frequencies, the skin effect and proximity effect will produce a complex distribution of currents in each conductor as well as induced currents in other nearby conductors that cannot be calculated with simple analytic expressions as can be done at dc. In each simulation we used 91 μm rf electrodes, a 70 μm centre electrode, 10 μm gaps around the rf electrodes and 5 μm gaps elsewhere. In each simulation the microwave power input onto each electrode was adjusted to null $B_{\mu w}$ at the ion whilst giving $B'_{\mu w}=10\text{ T/m}$. Figure 7.4 shows the resulting surface current distributions in each case. Table 7.1 gives a summary of the performance of each layout. The fig-

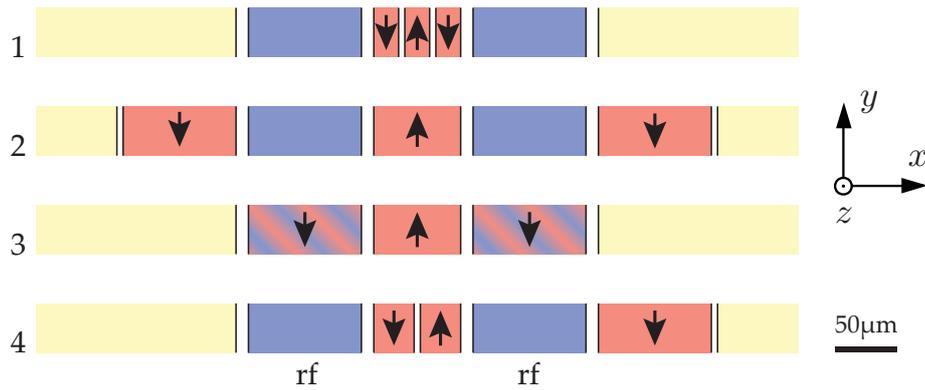


Figure 7.3: The four different microwave conductor layouts that we consider. The conductors extend ~ 2 mm in the y direction. In each the rf electrodes (blue) are in the same place. Microwave electrodes are coloured red and the ground plane is yellow. In layout 3 the rf and microwave signals are combined on the same electrodes (blue/red stripes). The relative sense of the microwave currents is indicated by the arrows.

ures of merit used are the peak surface current density on the trap $J_{surf,peak}$, the maximum average volume current density in a single electrode $J_{ave,max}$ and the total current required. The current densities are important as the power dissipated in the electrodes is proportional to the square of current (see section 7.1.5.4 for further discussion of power dissipation issues). To give an indication of what is realisable, a gold wire on sapphire of similar size has been measured to support a dc current density of up to $\sim 6 \times 10^{10}$ A/m² [GKW⁺04]. From this table layout 3 gives the lowest current densities at these dimensions and is therefore the one that we pursued. It does however add the complication of having to combine the rf and the microwaves on the same electrodes. A more exhaustive investigation could be carried out by varying the different electrode dimensions for every layout (within the constraints of figure 7.1); however we decided the advantages of layout 3 were clear enough to proceed directly to the optimisation of that design.

7.1.4 Impedance matching and build-up resonator

In order for the microwave power to pass through the trap, and not be reflected from it, the trap should be impedance-matched to the 50Ω drive electronics.

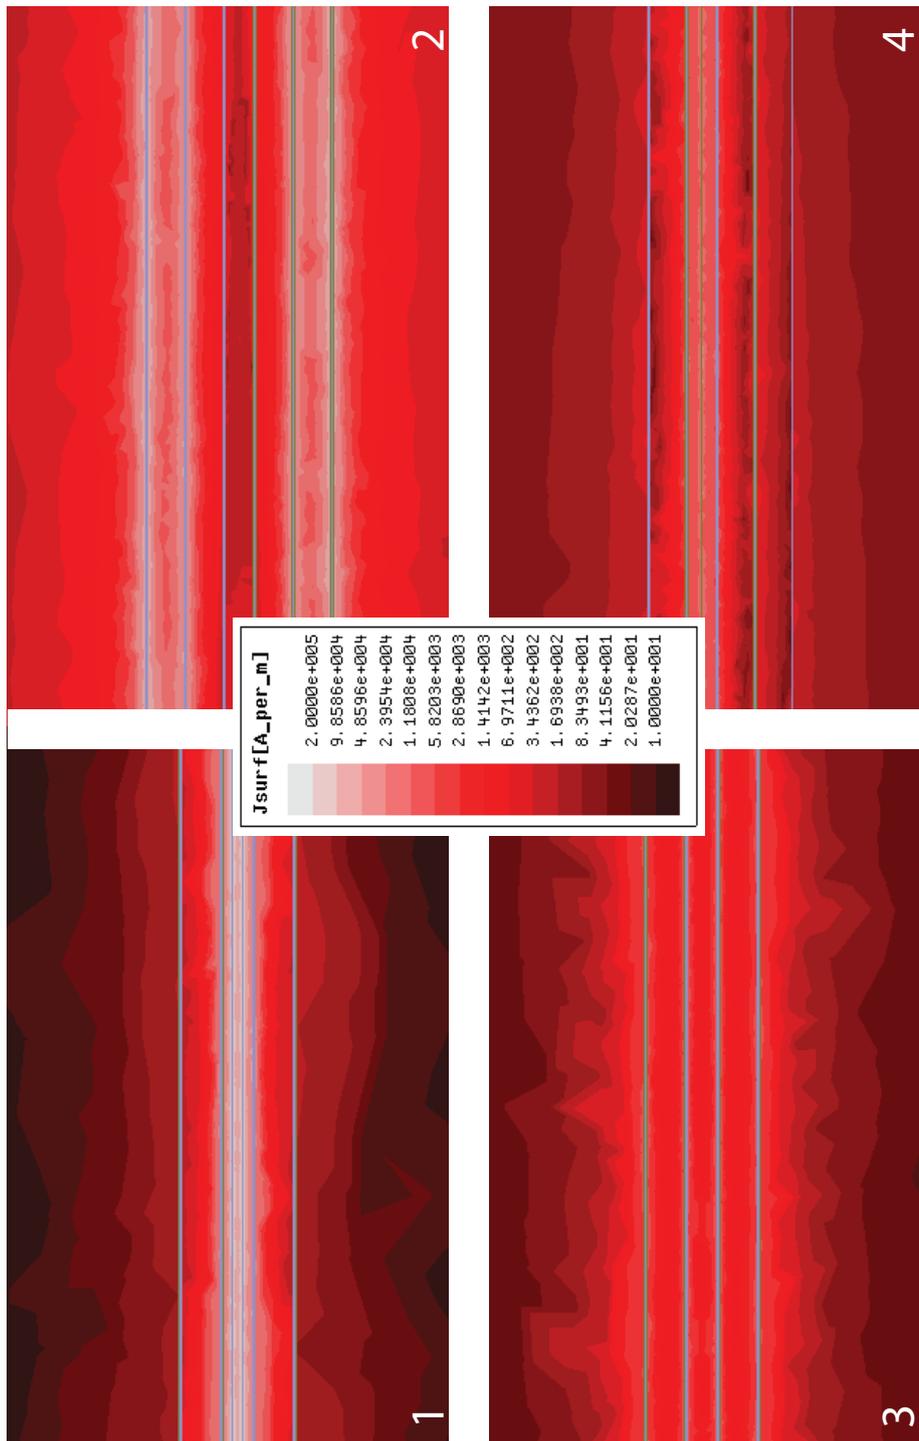


Figure 7.4: The surface current densities in each of the layouts when the three electrodes are driven so as to give $\mathbf{B}_{\mu w} = 0$ and $B'_{\mu w} = 10\text{T/m}$ at the ion.

Table 7.1: Simulation results for each layout of figure 7.3 when the microwave currents are adjusted to null $\mathbf{B}_{\mu w}$ at the ion and give $B'_{\mu w}=10$ T/m. $J_{surf,peak}$ is the peak surface current density on the trap and $J_{ave,max}$ is the maximum average volume current density in a single electrode.

Layout	Total current (A)	$J_{surf,peak}$ ($\times 10^3$ A/m)	$J_{ave,max}$ ($\times 10^9$ A/m ²)
1	15.7	250	62
2	16.7	100	10
3	2.8	16	1.9
4	2.9	36	6.7

Once the microwaves have passed through the trap region they can be terminated by a matched resistive load into which they will be dissipated [Tre08]. The microwave power in the trap could be increased by following the trap with a short or open circuit. This would work by reflecting the microwaves, causing them to pass through the trap a second time. This would create a standing wave so care would have to be taken to place the trap at a current anti-node. This occurs at $l = n\lambda/2$ away from a short and $l = (n + 1/2)\lambda/2$ away from an open circuit, where n is an integer. This technique could be extended yet further by placing the trapping region in a resonant structure that would allow many passes of the microwaves. Such a technique has been used by the superconducting qubit community to enhance (electric field) coupling between microwaves and transmon qubits [MCG⁺07], though these experiments are in the very different regime of single microwave photons.

7.1.4.1 Half-wave resonator

From the investigation in section 7.1.3.1 we found that it is advantageous to use the rf electrodes as two of the microwave conductors. This precludes us from shorting these electrodes to ground and so our resonator must be open circuit beyond the trap. This means the smallest resonant structure we can use is a half-wave resonator with the trap in the centre (where the current anti-node lies).

7.1.4.2 Coupling elements

An open-circuit $\lambda/2$ resonator would typically be coupled into using a series capacitor (often implemented as a simple gap in the transmission line) which, near resonance, makes the resonator look like a series RLC circuit. For optimal power transfer into the resonator this capacitor should be sized such that the resonator is critically coupled. This occurs when the impedance of the resonator on resonance is equal to Z_0 , the characteristic impedance of the feed line (typically $50\ \Omega$). A full analytical treatment for this system is given in Pozar [Poz05]. For our case however capacitive coupling is not possible as we also need to be able to apply dc and rf voltages to the electrodes. The capacitor can however be replaced with a $\lambda/4$ section of coplanar waveguide (CPW) of higher impedance than Z_0 which is electrically equivalent [ZSM05] to a capacitor.

7.1.5 Technical implementation

7.1.5.1 Trap substrate

Before we proceed further with design optimisation we need to fix our choice of substrate. This is because the dimensions of the microwave structures will depend on the relative permittivity η_r of the substrate. The surface trap in chapter 3 was fabricated on a crystalline quartz substrate. As the trap electrodes will be carrying, and dissipating power from, large microwave currents we require a substrate with a better thermal conductivity than quartz. Commonly used insulating microwave substrates are listed in table 7.2 (semiconductors are excluded because of their high microwave loss). We chose sapphire as it has the highest relative permittivity, which will make the resonator as small as possible, and relatively high thermal conductivity. CVD diamond and BeO have the best thermal properties but diamond is hard to work and expensive and BeO is toxic. AlN was rejected for our initial trap as it is a polycrystalline ceramic and we were unsure whether it would be compatible with our fabrication process. It

has since been used for an ion trap substrate [OWC⁺11] and therefore is worth future consideration if thermal conductivity becomes a limiting factor.

Table 7.2: Properties of commonly used microwave substrates. Modified from [CBB⁺05]. Sapphire is a uniaxial crystal with a relative permittivity of 11.53 parallel to the c-axis and 9.35 perpendicular to it.

Substrate	Thermal conductivity ($\text{Wm}^{-1}\text{K}^{-1}$)	Relative permittivity η_r	Loss tangent ($\tan \delta$)
AlN	180	8.5	3×10^{-4}
Alumina	30	9.8	1×10^{-4}
BeO	250	7	8×10^{-4}
BN	28	4.1	1×10^{-4}
CVD Diamond	800-1800	5.7	1×10^{-4}
Fused Silica	1	3.9	1×10^{-4}
Quartz	7	4.5	1×10^{-4}
Sapphire	45	9.4/11.5	1×10^{-4}

Sapphire is a uniaxial crystal with a relative permittivity of 11.53 parallel to the c-axis and 9.35 perpendicular to it. The c-cut is typically used for microwave substrates otherwise the properties of the circuit structures would depend on their orientation on the substrate, unnecessarily complicating the design. The loss tangent is 8.61×10^{-5} parallel to the c-axis and 3.01×10^{-5} perpendicular to it at 10 GHz.

We cut away the unused sections of the substrate on either side of the ground plane, to prevent the laser beams clipping the substrate. Without these cutouts the substrate would be 30 mm wide and the $1/e$ radius of a 397 nm beam at the substrate edge would be $> d = 75 \mu\text{m}$ (for any beam parameters) so clipping would be unavoidable. This could have been done instead by using two substrates, one with the 50Ω feed lines on and one with the resonator and trap on; however this would have added the complication of a wire bond transition between the two substrates. Instead we cut 2" circular wafers¹ to shape by gluing them together, five at a time, in a stack and then cutting the shape out with a

¹MTI Corp. single crystal Al_2O_3 , 0.5 mm thick, 10 Angstrom average surface roughness

diamond wire saw. This technique does not produce a straight cut so the back edge of the stack was ground flat to ensure that the substrate lies flush against the coaxial connections (see section 8.1.1.2).

7.1.5.2 Design optimisation

All the simulations in this chapter were carried out using Ansys HFSS. HFSS uses the finite element method to simulate electromagnetic structures at high frequencies. The advantage of this method over circuit model and 2D electromagnetic simulators (such as will be used in section 7.2.5) is that arbitrary 3D structures can be simulated and the resulting fields can be analysed at any point. The drawback is that it is time-consuming to set up and to verify HFSS models, and to compute the solutions to them (our trap takes about 30 mins to solve on a high-end workstation).

Figure 7.5 shows the final optimised design. As well as the basic elements discussed in the previous subsections there are a few other notable design features. The thickness of the electrodes was chosen to be $5\ \mu\text{m}$ because the skin depth at 3.2 GHz is $1.4\ \mu\text{m}$ so thicker electrodes would not reduce loss but would be harder to fabricate. The dc electrodes are set back from the rf electrodes with a thin section of ground plane in between. This is because otherwise the gaps between the dc electrodes would prevent current flowing in the ground plane, causing a reflection of microwave power. The microwave electrodes are also flared out away from the trap region; this is to reduce losses and spread out the thermal load.

Two main strategies were adopted in the optimisation of the design. The first involved systematically changing the electrode dimensions in the trap region and calculating the current densities required to produce a given $B'_{\mu w}$ at each point. The second involved varying the resonator and coupling element dimensions to obtain the best coupling of microwaves into the trap region. This

optimisation made use of the automated optimisation algorithms included in HFSS as simple optimisation goals can be expressed in terms of S-parameters (see section 8.2.1). The parameters that were optimised were the widths of the centre and rf trapping electrodes, the dc electrode set-back, the length of the resonator, the length and width of the coupling elements and the width of the gap at the open end of the resonator.

7.1.5.3 Results

Using the final optimised design, we calculated the microwave power required on each electrode to null $\mathbf{B}_{\mu w}$ at the ion whilst giving a gradient of $B'_{\mu w}=10\text{ T/m}$. Figure 6.12 shows the simulated behaviour of $\mathbf{B}_{\mu w}$ in the region around the ion. This gradient is large enough to drive a gate in 0.93 ms (see 6.6.3). The required powers were 624 mW on the centre electrode and 854 mW on the rf/ μw electrodes (4.66W in total with both sidebands). The resulting currents are 643 mA in the centre electrode, 839 mA in each rf/ μw electrode and 427 mA in each ground plane (the section between the rf/ μw electrodes and the dc electrodes). Figure 7.6 shows the surface current distribution on the trap. At these currents the field at the ion from the centre electrode alone is $\mathbf{B}_{\mu w} = 499\ \mu\text{T}$ (parallel to the substrate) and from one rf/ μw electrode alone it is $\mathbf{B}_{\mu w} = 586\ \mu\text{T}$ (at 65° to the substrate).

7.1.5.4 Thermal effects

A crude estimate of the temperature differential between the electrodes and the substrate can be obtained by assuming that half the microwave power is dissipated in the trap region. This region has an electrode surface area of $2\text{ mm}\times 402\ \mu\text{m}$ and the thermal conductivity of a gold-sapphire interface has been experimentally measured as $3.5 \times 10^6\text{ WK}^{-1}\text{m}^{-2}$ in similar atom chip experiments [GKW⁺04]. This implies $\Delta T \approx 0.8\text{ K}$ when the trap is driven

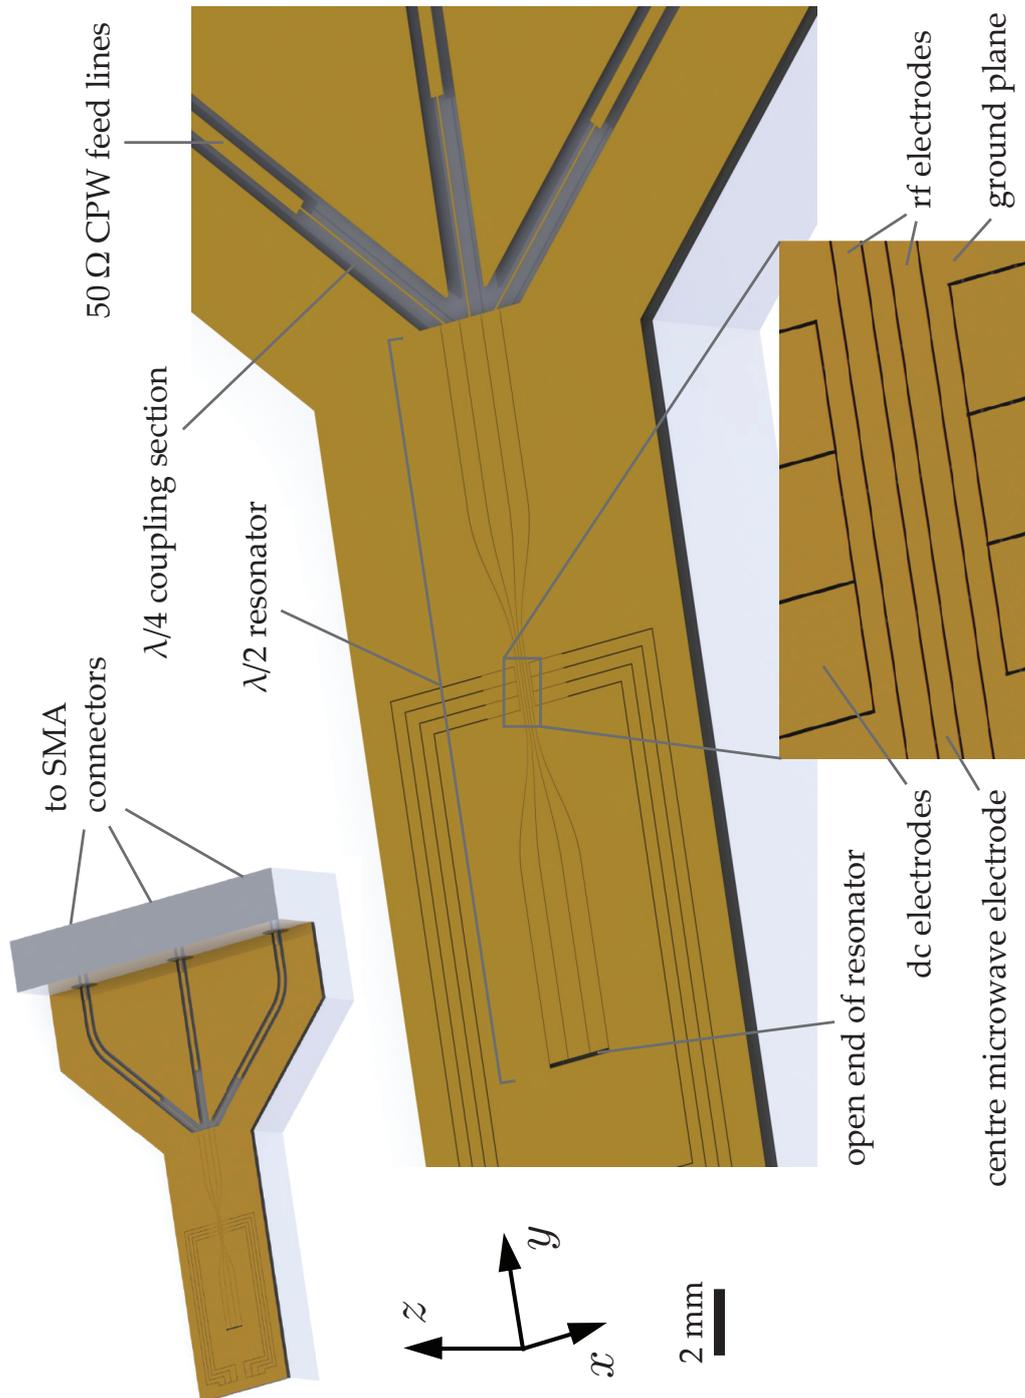


Figure 7.5: Electrode layout of final optimised trap. The upper inset shows the whole substrate and the lower inset shows the trap region. The microwaves are fed in on the three coplanar waveguide (CPW) feed lines at the right. Rf and dc are also combined on the same path.

7. DESIGN OF TRAP AND ELECTRONICS FOR MICROWAVE-DRIVEN QUANTUM LOGIC

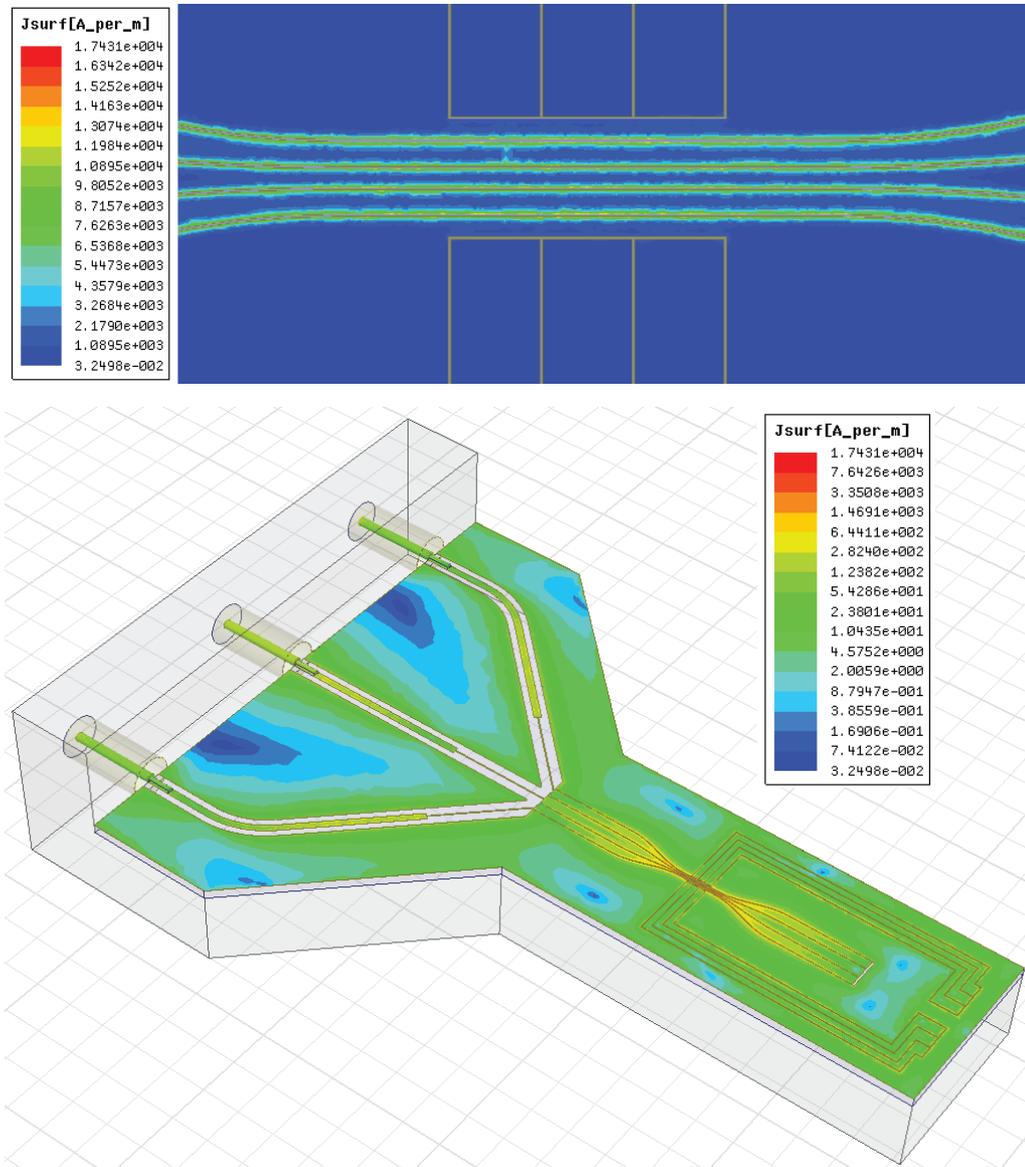


Figure 7.6: Surface current density plot of the central trapping region plotted with a linear scale (top). The current crowding at the edges of the electrodes due to the proximity effect is clearly visible. Below is the same plot but showing the whole trap using a log scale.

with 4.66 W. A further crude estimate of the substrate heating can be obtained by assuming all the power from a single microwave pulse goes into the $2\text{ mm} \times 0.5\text{ mm} \times 0.5\text{ mm}$ volume of substrate directly underneath the trap region. From the heat capacity of sapphire and assuming the pulse is applied for 1 ms the local substrate heating from a single pulse will be $\sim 2\text{ K}$. The slow heating due to many pulses is somewhat more complicated to simulate but if it was required then the HFSS simulation could be linked to ANSYS Multiphysics which would simulate the thermal behaviour.

7.2 Microwave drive system design

This section deals with the design of the microwave electronics. In order to drive the entangling gate the electronics must drive a pulse of microwaves into each electrode with the form

$$A_r \cos(\omega_r t + \phi_r) + A_b \cos(\omega_b t + \phi_b) \quad (7.1)$$

where ω_r and ω_b are the red and blue sideband frequencies. The amplitude and phase factors need to be separately controllable for each electrode. The electronics must also be able to drive pulses into a single electrode at the required carrier transition frequencies for single qubit operations, initialisation and readout.

Figure 7.7 gives the layout of the system which generates the microwave pulses required for qubit initialisation and quantum logic gates. All the parts, except the synthesisers, are connectorised packages with SMA connectors and are listed in table 7.3. Using connectorised components increases the bulk of the system considerably and adds connectorisation losses but allows much greater flexibility which is important as modifications are likely as work progresses toward the gate. The drive system up to the power amplifiers could in future be constructed on a single PCB using surface-mount components which would decrease the system size significantly. All the components are rigidly attached to

the optical table and connected via I-pieces or short lengths of semi-rigid SMA cable². This is done to reduce losses and ensure high phase and amplitude stability (see [RG02] for an example of this type of construction in demanding phase measurements). The synthesisers sit on a shelf above the optical table and are connected via test-quality flexible cables³. SMA connections are all tightened to 8 in.-lbs with a torque wrench to ensure reliable connections. All the amplifiers and terminations in the system are passively heat sunk to avoid the vibrations which would be caused by fans. In the case of the main power amplifiers two large heat-pipe radiators, taken from commercial CPU coolers, were also added as they can dissipate 75 W of heat. Synthesiser levels and frequencies are controlled over GPIB. All the synthesisers are externally referenced to the 10 MHz output of a rubidium frequency standard⁴.

The system consists of six main blocks (figure 7.7): carrier drive, initialisation drive, two sideband drives, power amplification and rf/microwave diplexer. Each of these blocks is now described.

7.2.1 Carrier drive

Carrier pulses, at the qubit frequency, are used to drive single-qubit rotations (see section 6.4.2). A synthesiser (p in figure 7.7) generates microwaves at the qubit frequency which is then split into two with a power divider (c). The first channel is used to drive rotations about the x-axis of the Bloch sphere and the second is used to drive rotations around an arbitrary axis in the x-y plane. This is achieved by using an IQ mixer (h) to vary the phase of the second channel whilst keeping the amplitude the same as in the first channel (an attenuator (q) is put in the first channel to cancel the insertion loss of the mixer). Each channel has a high isolation switch and is combined with the initialisation drive on a power

²RG402C, 95.1 pF/m, 0.69 dB/m loss at 3 GHz

³Minicircuits CBL series

⁴Stanford Instruments FS725 with Allan variance of $< 2 \times 10^{-11}$ over 1 s

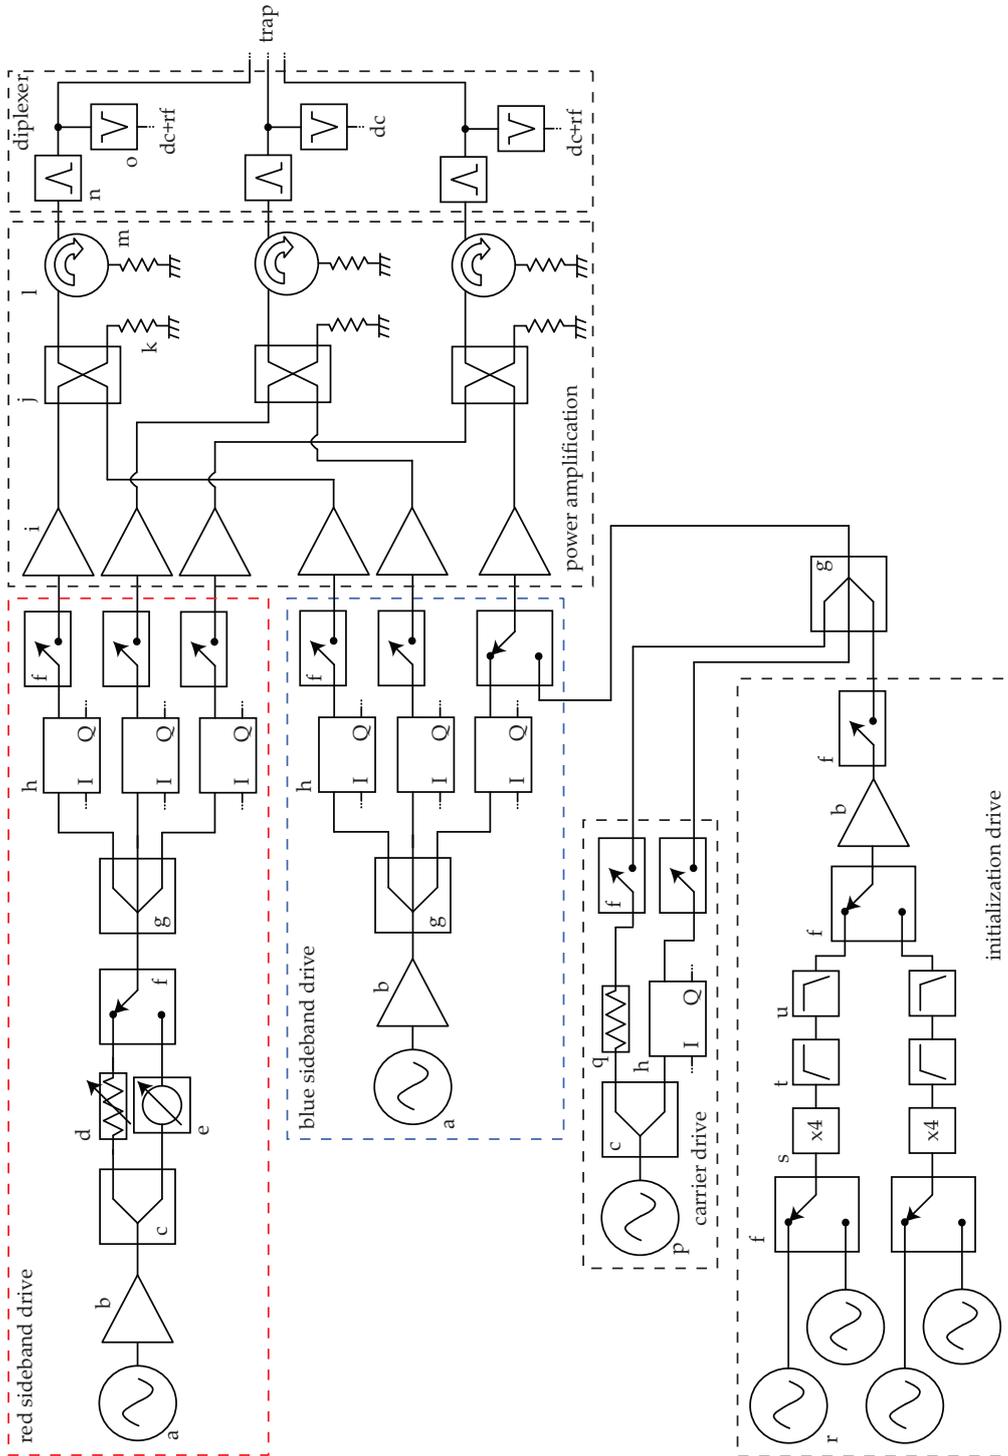


Figure 7.7: Schematic of the microwave drive system.

7. DESIGN OF TRAP AND ELECTRONICS FOR MICROWAVE-DRIVEN QUANTUM LOGIC

Table 7.3: List of parts in figure 7.7 and some nominal specifications that are particularly relevant for their application here. The insertion loss (IL) for splitters and hybrids is the loss above the theoretical loss for a perfect splitter. Some acronyms used in this table: MCL = Mini-Circuits Ltd, IL = insertion loss, VSWR = voltage standing wave ratio, NF = noise figure and P_{1dB} is the 1 dB compression point.

	Description	Part	Relevant specifications
a	Synthesiser	Agilent E4426B	250 kHz-4 GHz, $P_{out,max}=+7$ dBm 0.01 Hz/0.02 dB resolution < -95 dBc phase noise at 1 kHz
b	Low noise amplifier	MCL ZRL-3500	16.0 dB gain, $P_{1dB}=23.9$ dBm NF=3.3 dB
c	2-way 0° power splitter	MCL ZAPD-4	0.4 dB IL, <0.4 dB amp. unbalance
d	Variable atten.	ATM AF064-10	<0.5 dB IL, 10 dB atten. range
e	Phase shifter	ATM P1404	<0.5 dB IL, 0-96° shift
f	Absorptive SPDT switch	MCL ZASWA-2-50DR	2.6 dB IL, 69 dB in-out isol. $t_{switch}=10$ ns, $t_{rise}=5$ ns
g	3-way 0° Power splitter	MCL ZA3PD-4	>0.47 dB IL 0.46 dB amp. unbalance
h	IQ mixer	Eclipse Microwave IQ2040MP4	LO $P_{in}=+10$ dBm, 5.5 dB conv. loss 2.0° phase & 0.2 dB amp. match DC-500 MHz IF freq. response
i	Power amp	MCL ZHL-16W-43	46.0 dB gain, $P_{1dB}=42.8$ dBm NF=5.8 dB, $VSWR_{out}=1.23$
j	90° Hybrid	Narda 4033C	0.2 dB IL, 22 dB isolation
k	Termination	ATM T1515	$VSWR<1.2$, $P_{in,max}=15$ W
l	Circulator	ATM ATC3-6	21 dB isolation, 0.35 dB IL $VSWR<1.25$
m	Termination	Weinschel 24-30-33 + MCL KARN-50+	30 dB atten. + 50 Ω term. $VSWR<1.05$, $P_{in,max}=50$ W
n	Bandpass filter	Custom	See section 7.2.5
o	Bandstop filter	Custom	See section 7.2.5
p	Synthesiser	Agilent E4422B	250 kHz-4 GHz, $P_{out,max}=+7$ dBm 0.01 Hz/0.02 dB resolution
q	Attenuator BW-S12W2+	MCL	12 dB attenuator 0.005 dB/°C tempco
r	Synthesiser	HP 8656B	0.1-990 MHz, $P_{out,max}=+13$ dBm 10 Hz/0.1 dB resolution
s	Frequency multiplier	MCL ZX90-2-11	$P_{in,800\text{ MHz}}=+10$ dBm $P_{out,3200\text{ MHz}}=-19$ dBm
t	Highpass filter	MCL VHF-2700A or MCL VHF3100	$f_{co}=2700$ MHz $f_{co}=3100$ MHz
u	Lowpass filter	MCL VLF-3000 or MCL VLF-3400	$f_{co}=3600$ MHz $f_{co}=3950$ MHz

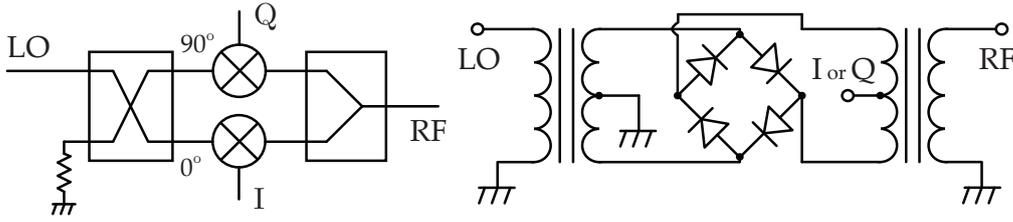


Figure 7.8: Schematic of an IQ mixer (left) and a schematic of one of the double balanced mixers inside the IQ mixer (right).

combiner (g) before being sent to the power-amplification stage. The carrier is applied to one of the outer electrodes as these give a $B_{\mu w}$ -field $\approx 65^\circ$ to the static B-field (see section 7.1.5.3) and thus can drive transitions of any polarization.

7.2.1.1 The IQ mixer

In order to control both the phase and amplitude of each microwave channel we employ IQ mixers. Figure 7.8 shows the internal layout of this device. The 90° hybrid splits the LO input into two arms, one 90° out of phase with the other. A double balanced mixer in each arm is then used as a modulator before the two arms are combined on a 0° power combiner. The output is thus described as

$$A_I \cos(\omega t) + A_Q \sin(\omega t) = A \cos(\omega t - \phi) \quad (7.2)$$

where $A = \sqrt{A_I^2 + A_Q^2}$ and $\phi = \tan^{-1}(A_Q/A_I)$. A_I and A_Q are set by the current driven into the I and Q inputs (see fig 8.13). Whilst we will initially drive the I and Q inputs with dc, they have a frequency response of up to 500 MHz and so could be driven by high speed arbitrary waveform generators (AWGs) in order to perform fast phase and amplitude pulse-shaping during the gate [Bia11].

The IQ mixers are driven by the same DAC as used in section 4.2.2, though an output buffer board is required as the DACs cannot supply the required current. This buffer is a lowpass (~ 30 Hz cutoff) op-amp filter with a gain of 0.15 and a 50Ω series resistor in the output to limit the current through the mixer diodes. This resistor and the gain resistor are all ± 15 ppm/ $^\circ\text{C}$ resistors for stability.

7.2.2 Initialisation drive

The four frequencies required to initialise the qubit are generated by four separate frequency quadrupled rf synthesisers (r). We use these rf synthesisers rather than microwave synthesisers as we do not require such a high specification and they are much cheaper. The frequency and amplitude of these pulses is set digitally (over GPIB) on the synth. An alternative would be to use a DDS system which would be less bulky but would not offer any performance improvements when using simple π pulses. DDS's inherent phase agility would allow for composite pulse schemes to be realised though and so may be implemented in the future. Quadrupling is done using single frequency doublers (s) with high pass filters (t) to block the second harmonics and low pass filters (u) to block the 6th harmonics. The quadrupling has low efficiency (~ -30 dB) so a pre-amp (b) is required afterwards. Adiabatic transfer sweeps could be implemented by using two AWGs, one connected to the FM input of all the synths and one connected to a mixer after the pre-amp to modulate the amplitude.

7.2.3 Sideband drive

Two synthesisers (a) generate microwaves at the ions' motional sideband frequencies, $\omega_{r,b} = \omega_0 \pm \omega_j \mp \delta$. Each of the synth outputs is pre-amplified (b) (as the IQ mixers require a fixed +10 dB LO input) before being split into three channels. Each channel has its own IQ mixer (h) and switch (f) which it passes through before the power amplification stage.

7.2.3.1 Phase flipper

In order to perform 180° phase shifts in one of the sideband drives (see section 6.6.3.2) we use a power splitter to split the output from one of the synths into two arms. In one arm we place a manual variable phase-shifter (e) and on the other a manual variable attenuator (d). A DPST switch is then used to

switch between the two arms, which it can do in ~ 10 ns. The attenuator and phase shifter are then adjusted so that the output of the two arms after the switch has the same amplitude but a 180° phase shift (this would not be required if we were driving the IQ mixers with AWGs as we could use these to shift the phase rapidly).

7.2.4 Power amplification

Each sideband drive channel has its own 16 W solid-state power amplifier (i). In principle the red and blue sideband could be combined before the power amplifier, which would save three amplifiers. However this would greatly complicate nulling the carrier for gate operation as, due to compression, the gain of the amplifier would depend on the amplitude of the other sideband. The carrier and initialisation drives can however share an amplifier with one of the channels though as they are not used at the same time as the sidebands. The pairs of amplifiers are combined using 90° hybrids (j) rather than 0° Wilkinson power combiners. This is because, when combining non-coherent signals, half the power will be dissipated in the internal resistor of the combiners. These resistors are generally small surface mount components and could easily be destroyed or at the very least could give rise to thermal transients in the device's properties during a pulse. Hybrids on the other hand have two output ports where with half the input power appearing at each and little internal dissipation (see figure 8.14). One of those ports can then be terminated with a high power termination (k). After the combiner is a circulator (l) which allows forward-going power to the trap but diverts any power reflected from the trap to a high-power terminator (m). This isolates the power amplifier from any impedance mismatches with the trap. This is a useful technique for eliminating load-pull which occurs when reflections from a mismatched load cause changes in the power output of the amplifier.

7.2.5 RF/microwave diplexer

The design of the trap requires that each of the three central trap electrodes be connected to a diplexer to allow the combination of a microwave channel, an rf channel and a dc bias. This is done using four separate filters labelled a-d in the diplexer block diagram of figure 7.9. The centre electrode has an identical circuit but with its rf input connected to ground to prevent rf pickup on this electrode. The diplexer circuit was designed using the circuit simulator in AWR Microwave Office (MWO)⁵. This simulator allows one to wire together elements that can either be discrete components such as capacitors or structures patterned on the circuit board such as transmission lines. These elements contain a model for their electrical behaviour based on closed-form equations, electromagnetic simulation or vendor-provided data (in the case of discrete components). From the circuit the software can calculate various outputs such as port scattering parameters. Optimisation algorithms can then be used to vary selected elements (e.g. length of a transmission line or value of capacitor) to try and obtain the required output parameters. In principle HFSS could also be used to design the diplexer but for simple filters a circuit model is much easier to work with and faster to optimise. It is also difficult to incorporate discrete components into HFSS.

The ideal design should have minimal loss in the microwave path to maximise power at the trap and minimal capacitance in the rf path in order to reduce the load on the rf transformer. The diplexer is designed to be fabricated on high frequency ceramic-PTFE composite board⁶ for low loss and high temperature stability. The microwave traces on the circuit between the filters are 50 Ω microstrip, which on this substrate is 1.24 mm wide. Figure 7.9 shows the microwave office model and the final layout.

⁵this software was generously donated by AWR

⁶Rodgers Corp. RO3003, 0.50 mm thick, 17 μm copper clad, $\epsilon_r = 3.00$ (+13 ppm/ $^\circ\text{C}$), dissipation factor = 0.0013

7.2.5.1 Filter a

This filter must be open circuit for the dc, a small capacitance at rf and pass microwaves with a low insertion loss. We chose a coupled line bandpass filter which passes frequencies where the length of the lines is $\lambda/4 + n/2$. Our filter has three sets of lines to achieve tighter coupling which widens the bandwidth and decreases the insertion loss.

7.2.5.2 Filter b

This filter must have a dc path, a low inductance rf path, low capacitance to ground and it must block the microwaves. Open circuit $\lambda/4$ microstrip stubs act as a series LC circuit to ground and can be used to construct a microwave band-stop filter. A second-order filter was chosen as its insertion loss (> 33 dB from simulation) is already low enough to cause negligible microwave losses and a higher order filter would just add to its capacitance.

7.2.5.3 Filter c

An rf choke⁷ is used here to block the rf but pass the dc. The largest choke available with a resonant frequency above the rf frequency is chosen. This choice will present the highest impedance to the rf. It is also important this choke can handle the high voltage across it. There is a capacitor⁸ to ground after the choke to shunt any rf that leaks through the choke to ground and ensure it does not get to the DAC box. The choke and this capacitor also filter any noise picked up on the cables that carry the dc voltages from the DAC box.

⁷Epcos B82111EC25, 100 μ H, 500V, $f_{res}=55$ MHz

⁸Murata LLL, 0.47 μ F, 10 V

7.2.5.4 Filter d

As this only has to block dc and pass rf a simple capacitor⁹ suffices here. A capacitor with a low dielectric temperature coefficient is chosen for stability.

⁹Kemet C2225C103JCGAC, 10nF, 500V, C0G ceramic (± 30 ppm/ $^{\circ}$ C)

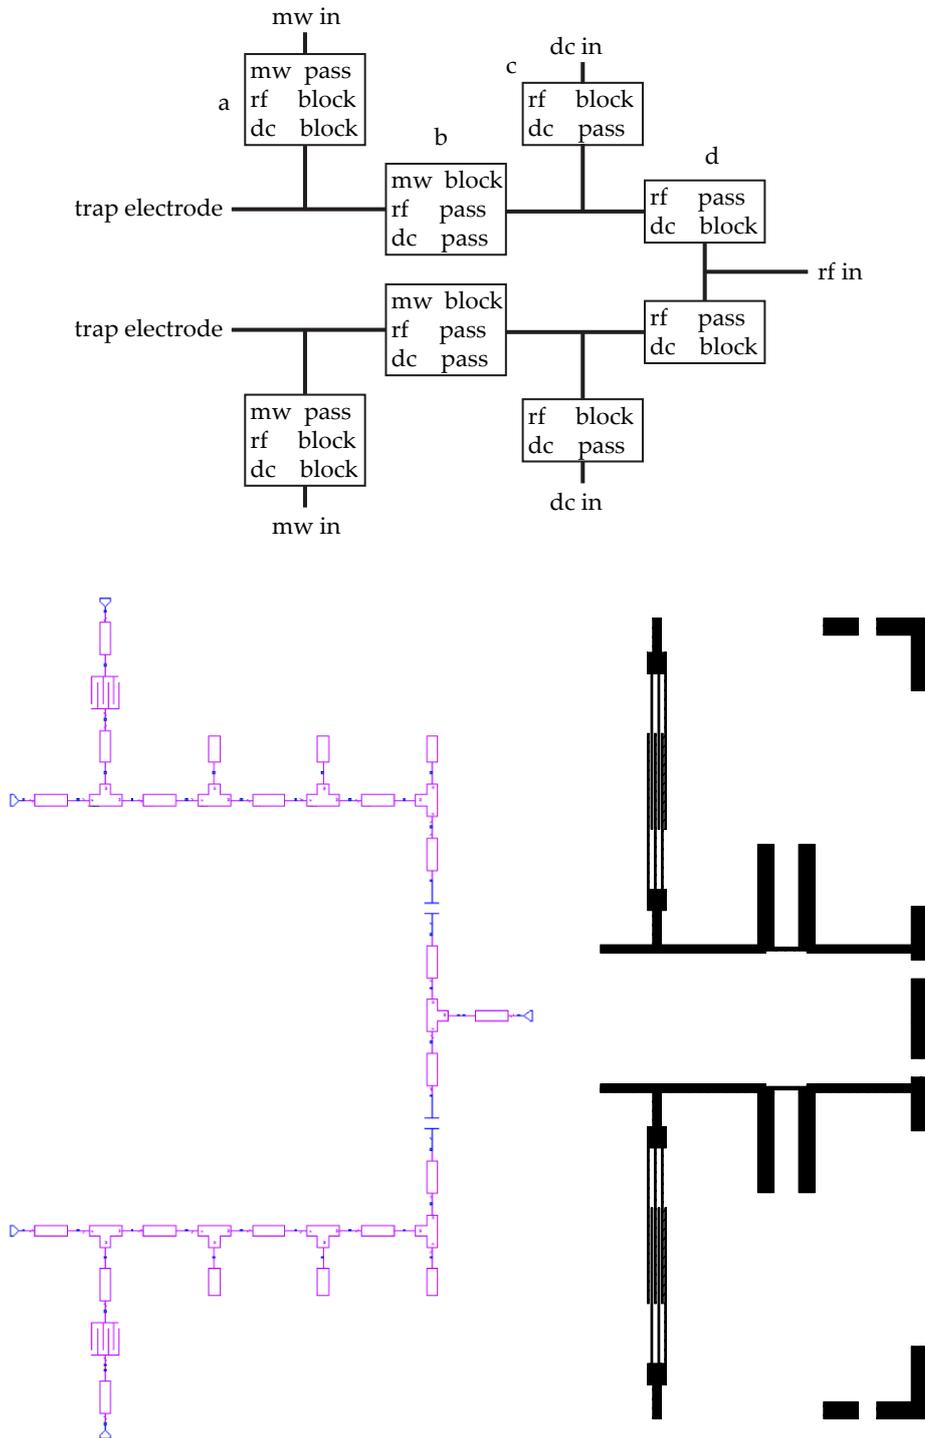


Figure 7.9: Block diagram of the diplexer (top), design of diplexer in Microwave Office (bottom left) and actual circuit layout (bottom right).

8

Construction and testing of microwave-driven quantum logic trap

I've got 50000 Watts of power.
I want to ionize the air.
This microphone turns sound into electricity.
Can you hear me now?
Out on Route 128, the dark and lonely.
I got my radio on.
Can you hear me now?
It's the end of radio.

Shellac

The first half of this chapter contains details of the apparatus constructed to realise the designs laid out in the previous chapter. As this experiment replaces the earlier experiments in this thesis, any apparatus not explicitly described can be assumed to be as described in chapter 2. The second half of the chapter details experimental results that have been obtained in characterisation of the trap and work toward the goal of microwave-driven two-qubit entanglement.

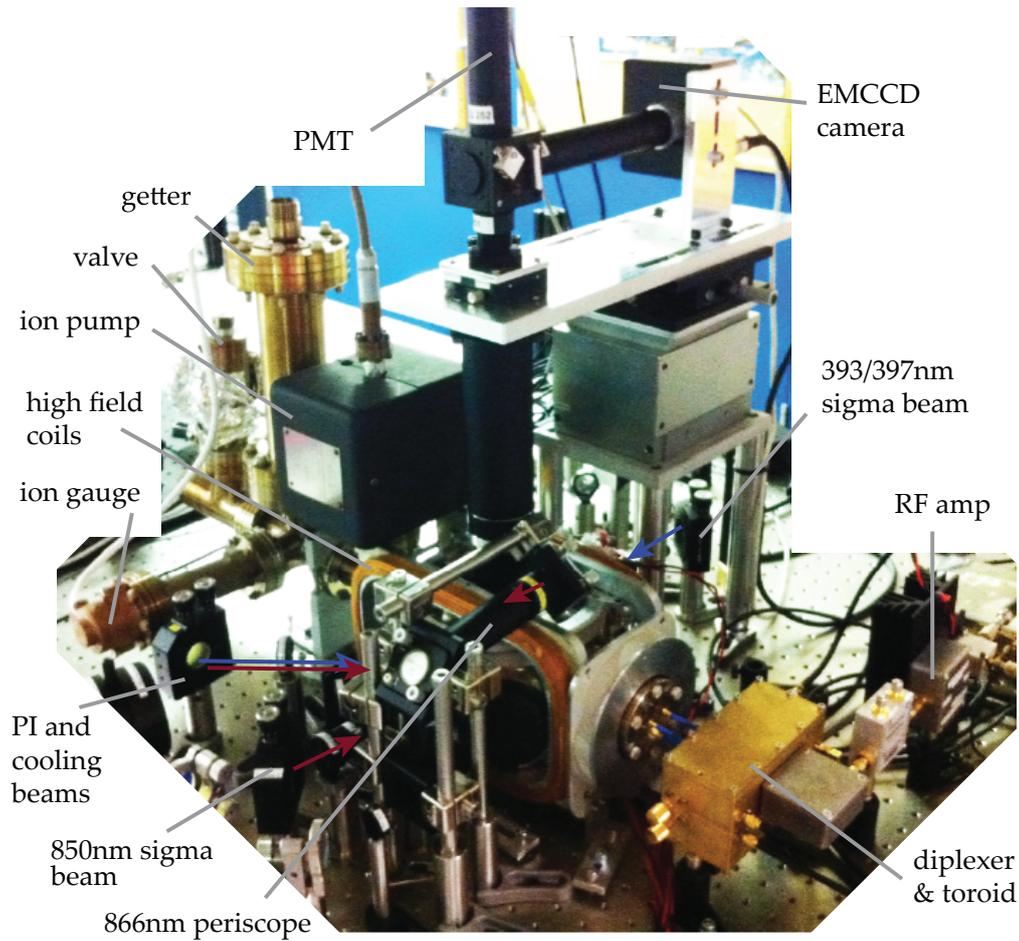


Figure 8.1: Photograph of vacuum system on the optical table with associated optics and electronics.

8.1 Experimental apparatus

Figure 8.1 shows the completed experimental apparatus described in detail below. Figure 8.2 shows the main octagon of the vacuum system schematically

8.1.1 Vacuum system and trap electrical connections

The vacuum system is built around a CF flange octagon¹. Above the trap, at a distance of ~ 40 mm, is a Kodial glass imaging window. It has an anti-reflection (AR) coating at 397 nm on the outside and a conductive indium tin oxide (ITO)

¹Kimball Physics MCF600-SphOct-F2C8, 316L stainless steel with 2x 6" CF flanges and 8x 2.75" CF flanges

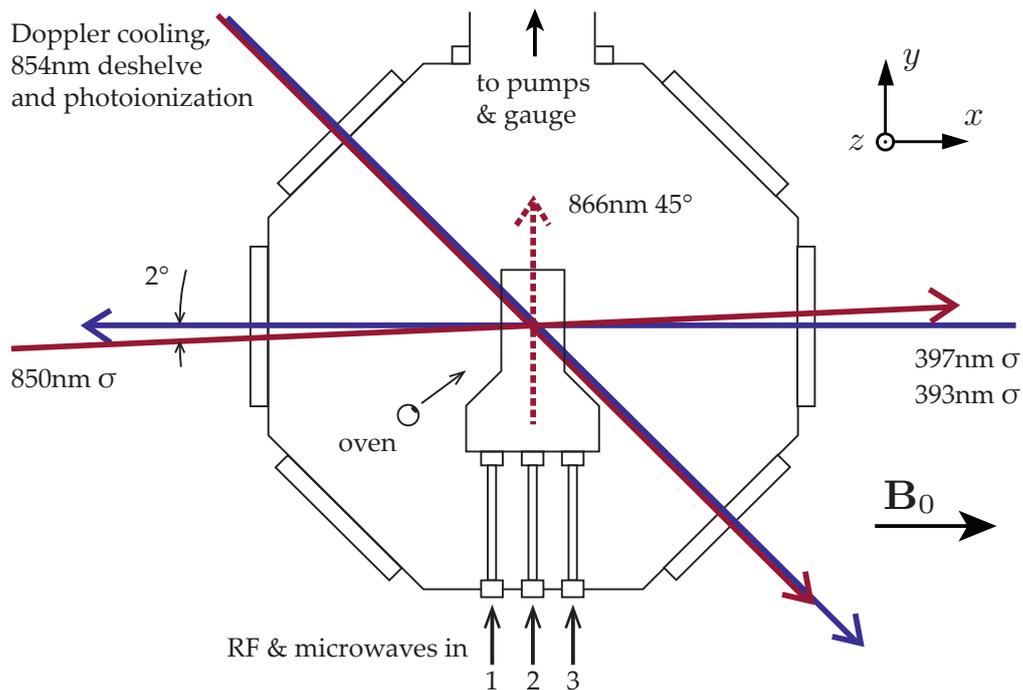


Figure 8.2: Schematic of the vacuum system octagon with the trap in the centre. The dashed 866 nm 45° beam comes through the top window and reflects off the trap before exiting back through the top window.

coating on the inside to prevent charging. Transmission at 397 nm is 76%. The side windows are fused silica with an AR coating for the UV lasers on both sides. The 6" flange on the bottom of the octagon is a custom flange² that contains the oven feedthroughs, the dc feedthrough and tapped holes to mount the trap base. The system is pumped using an ion pump³ and a non-evaporable getter⁴. The final bake of the system was performed at 190 °C (see appendix A for full details). With the oven off the pressure is below the measurement sensitivity of the ion gauge⁵; with the calcium oven on at typical loading current (5A) it is $\sim 1.5 \times 10^{-11}$ Torr.

²LewVac 6" CF flange with two 2-pin copper 15A power feedthroughs and a 25-way sub-D feedthrough

³Varian VacIon Plus 20 Diode, 271/s

⁴Saes Getters GP100 Mk5, ST707 alloy, 2401/s H₂ pumping speed

⁵Varian UHV-24P with x-ray limit of 5×10^{-12} Torr

8.1.1.1 Trap mount and dc connections

The trap is attached to an aluminium mounting block (see fig 8.3) with a thin layer of Epo-Tek 353ND epoxy. This ensures that the trap is well heat-sunk to a large thermal mass. It was hoped that the compliance of the epoxy would ensure that, despite the differing thermal expansion coefficients of aluminium and sapphire ($22.5 \times 10^{-6} \text{ K}^{-1}$ and $4.5 \times 10^{-6} \text{ K}^{-1}$ respectively), the trap would not crack during the 190°C UHV bake. Unfortunately the substrate did develop a hairline crack underneath the $\lambda/4$ coupling elements; however the gold film still appears to be intact, as indicated by the ion's correct response to dc voltages, rf and microwaves applied to all three electrodes. This mounting block also incorporates the microwave connectors (see section 8.1.1.2) and the dc electrode connections (see figure 8.3). The dc connections are made by wire bonding the trap to copper wires. The copper wires are Kapton-dipped wires⁶ that were inserted into a recess in the base and then glued in place with Epo-Tek 353ND epoxy. Once the epoxy had cured the top of the wires were filed down to expose the conductors and polished to provide a flat surface for wire bonding. The wire bonds are gold and do not adhere well to the copper so they are secured with conductive epoxy⁷. Single-layer capacitors to ensure good rf grounding (as used in section 3.2.1) are also wire-bonded to these wires. The capacitors are grounded by gluing them directly to the base with the conductive epoxy. The wires from the trap and the wires from the D-type connector are connected with barrel connectors. Unfortunately during the bake one of these barrel connectors sagged against another one and shorted them together. Even though this means two electrodes are constrained to be at the same voltage there are still enough degrees of freedom to generate the trapping potentials we require (although slightly higher voltages are required). This will be simple to correct next time the vacuum system is opened but the risk of damaging something else

⁶Kurt Lesker FTAK06010, 0.69 mm outer dia. 0.63 mm conductor dia.

⁷Epo-Tek H20E, silver-loaded heat-curing epoxy with volume resistivity of $\leq 0.0004 \Omega\text{-cm}$

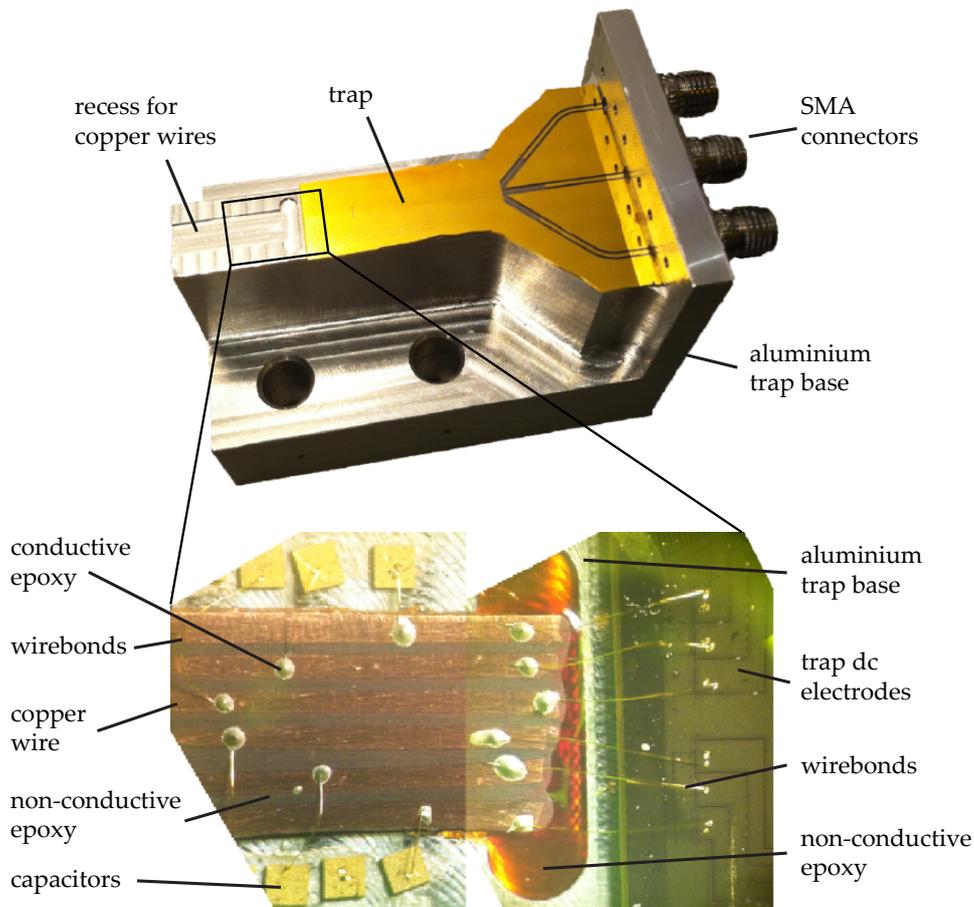


Figure 8.3: Photograph of the trap mounted on the aluminium base (top) and the connections between the trap dc electrodes and the copper wires which go to the dc feedthrough (bottom).

during another bakeout makes doing so now an unnecessary risk.

8.1.1.2 Microwave connections

A method was required to make the transition, in vacuum, from coaxial cable to the grounded coplanar waveguide on the trap. This was done by modifying a commercial 2.92mm bulkhead connector⁸ (see fig. 8.4). These connectors are compatible with SMA connectors but feature an air, rather than PTFE, dielectric. The stainless steel connector body and gold-plated beryllium copper contact are UHV compatible but the capture bead and dust cover are not. The capture bead

⁸Southwest Microwave 1012-16SF with 290-13P launch pin

8. CONSTRUCTION AND TESTING OF MICROWAVE-DRIVEN QUANTUM LOGIC TRAP

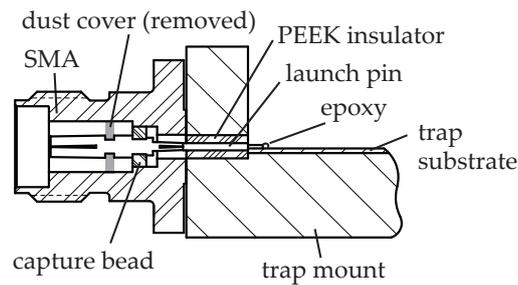


Figure 8.4: Cross-section of the SMA-to-trap transition (side-view).

is made from Ultem 1000 and was removed by dissolving it in dichloromethane (CH_2Cl_2). A new bead was made to the same dimensions using PEEK. The relative permittivity of PEEK and Ultem 1000 are 3.30 and 3.15 respectively so the impedance of the part will not be greatly affected by this change. The relative permittivity of the Kel-F dust cover is only 2.4 so we discarded it entirely (as replacing it with PEEK would not have been much better). There is then a launch pin through the aluminium trap base and onto the chip. This is supported by a PEEK insulator with an outer diameter of 2.3 mm, chosen to make it $50\ \Omega$. The insulator has a small groove along its length to prevent a trapped volume forming behind it. The launch pin was glued to the trap (see fig. 8.5) by conductive epoxy⁷ applied using a hypodermic needle on an xyz translation stage. At the same time the ground plane of the trap was also connected to the aluminium trap mount using the same epoxy. An in-vacuum SMA cable⁹ was attached to the connector and tightened to the specified torque (8 in.-lbs) *before* applying the epoxy, to prevent the epoxy connection being broken by the rotation of the pin.

Microwaves enter the vacuum system through SMA feedthroughs¹⁰ with an air-dielectric and a glass-ceramic seal which allows baking to 350°C . The feedthrough we used has a measured insertion loss of $<0.12\ \text{dB}$ and a return loss of $>26\ \text{dB}$ at 3.2 GHz. The in-vacuum SMA cables⁹ that connect the trap

⁹LewVac, 105 mm long, plugs are gold-plated brass plugs with gold-plated copper pin, PEEK insulator and assembled using UHV-compatible solder, cable is 26AWG kapton-insulated $50\ \Omega$ coaxial cable with silver-plated copper conductors and a capacitance of $95\ \text{pF/m}$

¹⁰Ceramtech 21363-01-CF, 3x double-ended SMA feedthroughs with grounded shields on a single 2.75" CF flange

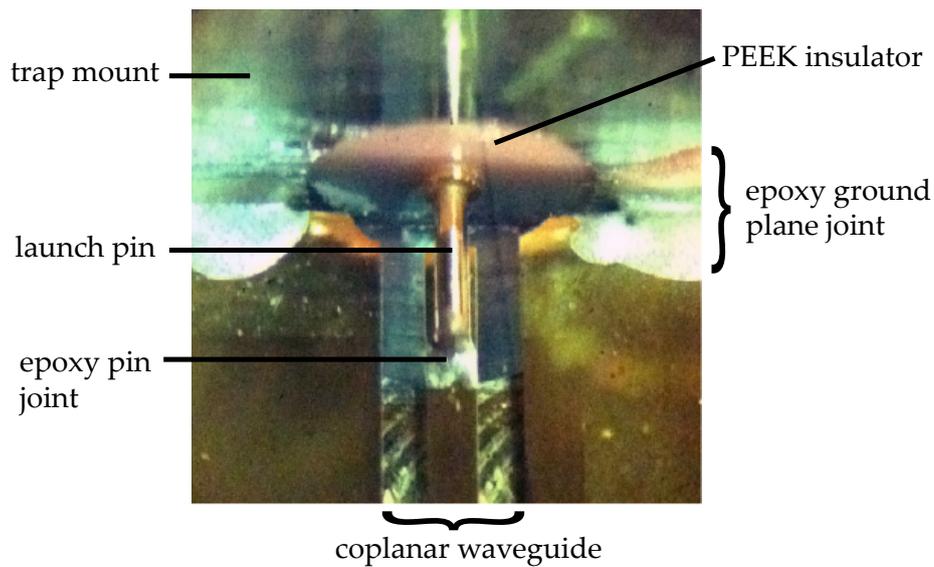


Figure 8.5: Photo of one of the launch pin exits and the epoxy which connects it to the waveguide transmission line on the trap. The aluminium trap mount is visible through the uncoated strips of the sapphire substrate.

to the feedthrough and have an insertion loss of <0.41 dB and a return loss of >22 dB. The connectors on this cable have had their PTFE dielectric replaced with PEEK by the supplier so are not quite $50\ \Omega$. However, without a better option available off-the-shelf, their performance was deemed good enough to make manufacturing our own cables unnecessary.

8.1.2 Imaging system

The imaging system used in this chapter differs from the one described in chapter 2. It is oriented vertically rather than horizontally, it has twice the magnification ($\times 16$) and it uses a different objective lens. This lens was custom designed for the Innsbruck group by Sill Optics¹¹. The numerical aperture of the imaging system is ≈ 0.3 and the overall detection efficiency is $0.23(2)\%$ with the slits closed up to give an effective aperture of $100\ \mu\text{m} \times 100\ \mu\text{m}$ (in the object plane). This efficiency is measured using a single trapped ion (by the method of

¹¹5-element lens, 40 mm diameter, 66.83 mm effective focal length, corrected for 6 mm fused silica vacuum window

[SKDB10]) and so includes fluorescence from the ion which is reflected from the trap's surface (see section 3.3).

8.1.3 Field coils

In order to generate large (~ 200 G) static magnetic fields at the ion we require water-cooled coils. We use a coil pair, each consisting of 16 turns of hollow copper wire¹², which together produce 1.0 G/A at the ion and have a combined resistance of 71 m Ω . The coils are driven by a pair of high current power supplies in series¹³ which can be controlled via GPIB. Cooling water flows along the hollow wires. A flow sensor¹⁴ in the water cooling circuit and thermal switches on the coil are used to interlock the power supplies and prevent overheating in the event of a cooling water failure.

8.1.4 RF supply

The trap rf voltage is generated by a synthesiser, amplified with a 800 mW amplifier¹⁵ and then stepped-up by a toroidal resonant transformer. The toroid (Micrometals T80-10) is mounted on a PTFE spacer inside an aluminium box bolted to the diplexer box (see figure 8.11). The resonant circuit has a Q of 65 and provides a voltage step-up of 9.6 at a resonant frequency of 38.7 MHz (deduced from the ion's secular frequency at a given rf input power).

8.1.5 DC supply

The dc voltages for the trap are generated by the same DAC and buffer system used for controlling the IQ mixers (see section 7.2.1.1). The dc connections for the rf electrodes and the centre control electrode are made via SMA cables to

¹²Wolverine Tube Europe OFHC copper 3.2×3.2 mm square with 2 mm dia. bore, kapton wrapped by S&W Wire Company

¹³Agilent 6671A, 8 V, 220 A

¹⁴Gems Sensors 155425BSPP

¹⁵Motorola CA2810C, 34.1 dB gain, noise figure < 4 dB

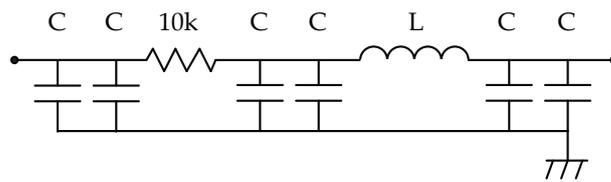


Figure 8.6: Circuit diagram of the dc filters. The inductor is a Coilcraft 1812CS ceramic chip with $L=22\ \mu\text{H}$ and an SRF of 40 MHz. The capacitors are Murata LLL series low ESL ceramic chips (25V) with $C=0.47\ \mu\text{F}$.

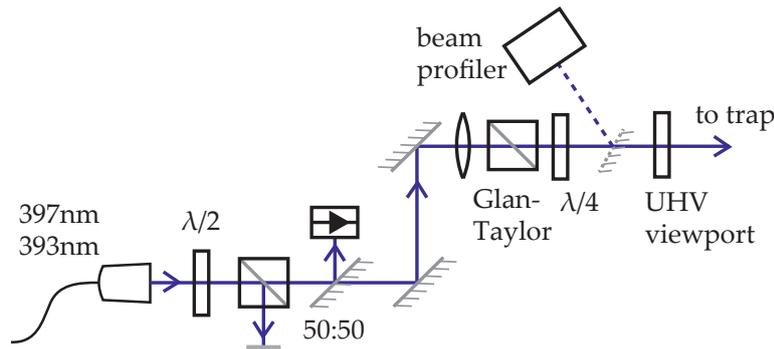


Figure 8.7: Optical layout for the UV sigma beam path. The mirror that directs the beam to the beam profiler is movable and only in place during alignment. The 850 nm sigma path has an identical arrangement.

the diplexer where they are filtered further by the capacitor and choke in the diplexer (see section 7.2.5). The dc electrodes are connected via a 6-way shielded cable with 15-way sub-D plugs. This cable connects to a filter box (see figure 8.6) mounted directly to the vacuum feedthrough.

8.1.6 Optical pumping beams

Figure 8.2 shows the sigma beam paths relative to the vacuum system and figure 8.7 shows the optical layout for each sigma beam path. High quality circular polarization is obtained by using a Glan-Taylor polariser¹⁶ and by mounting a zero-order quarter-wave plate on a kinematic mount so that it can be tilted about an axis normal to the beam to ensure the correct retardation. The quarter-wave plate can only be optimised for one wavelength however, so if the same

¹⁶Thorlabs GT-10A, 100,000:1 extinction ratio

optics are used at both 397 nm and 393 nm some error will be introduced by the 3.5 nm difference between them. If the waveplate is set at one of these two wavelengths then the polarization impurity (fraction of the wrong circular polarisation) introduced at the other wavelength is calculated to be 4.8×10^{-5} (or a 1.2×10^{-5} impurity at both wavelengths if the waveplate tilt is set in between). In practice the polarization error introduced by the UHV viewport is expected to be much greater than this, though we are yet to measure this on our vacuum system (an unmounted viewport was measured to introduce $\sim 10^{-4}$ polarization error). Techniques for producing higher quality viewports exist [SZW11]. The beam is aligned on the ion using the same technique as in section 2.1.4.2 and then the B-field direction is trimmed to ensure it is co-linear with this beam direction (by, for example, optimising optical pumping at 397 nm using $^{40}\text{Ca}^+$. The 850 nm sigma path has the same optical layout but comes in through the opposite viewport. This beam has a small angle ($\sim 2^\circ$) to the field (see section 6.5.0.2).

8.2 Microwave testing

In this section we will make extensive use of the scattering matrix $[S]$ to describe N-port microwave networks such as the trap. The elements of this matrix (the S-parameters) are defined as

$$S_{ij} = \frac{V_i^-}{V_j^+} \Big|_{V_k^+ = 0 \text{ for } k \neq j} \quad (8.1)$$

where V_j^+ is the amplitude of the voltage wave incident on port j and V_i^- is the amplitude of the reflected wave coming out of port i [Poz05]. We measured S-parameters using a Wiltron 37269A vector network analyser (VNA). The ports on the trap are labelled 1-3 (see figure 8.2).

8.2.1 Trap testing

Once the trap was installed in the vacuum system and baked out we measured the trap's S-parameters and compared them with the HFSS simulation (see figure 8.8). From these S-parameters the fraction of the power that is coupled into the vacuum system on port 1 is given by $1 - |S_{11}|^2 - |S_{12}|^2 - |S_{13}|^2$ (see figure 8.9 and table 8.1). The power coupled in is either dissipated in the trap, radiated or dissipated in the feedthrough and in-vacuum cabling. This feedthrough and cabling were measured separately before the system was assembled and their combined loss is also listed in table 8.1 and has been subtracted from the VNA data in figure 8.9. The Q of the microwave resonators is approximately 6 and matches the simulation well. The resonant frequency of the resonators is however around 5% higher than simulated. The possibility that this error is due to fabrication tolerances is investigated below and discounted. It seems likely though that this level of simulation accuracy is as high as is likely to be achieved without feedback from actual fabricated devices. In this manner it would be fairly trivial to revise the design with slightly longer microwave resonators. In any case though the modest Q of the resonators means we are still close enough to resonance that we couple most of our power into the resonators. There are also some smaller features on the plot that are not reproduced by the model but these may be due to resonances between the SMA cable, feedthrough and SMA connectors on the trap which are not included in the model.

Table 8.1: Fraction of incident power coupled into the vacuum system for each microwave input port at the three qubit frequencies we are considering. The cable loss is the loss of the feedthrough and in-vacuum SMA cable at each of these frequencies.

qubit freq (GHz)	port 1	port 2	port 3	cable loss
3.226	0.783	0.833	0.766	0.123
3.200	0.739	0.826	0.720	0.114
3.123	0.633	0.858	0.637	0.110

8. CONSTRUCTION AND TESTING OF MICROWAVE-DRIVEN QUANTUM LOGIC TRAP

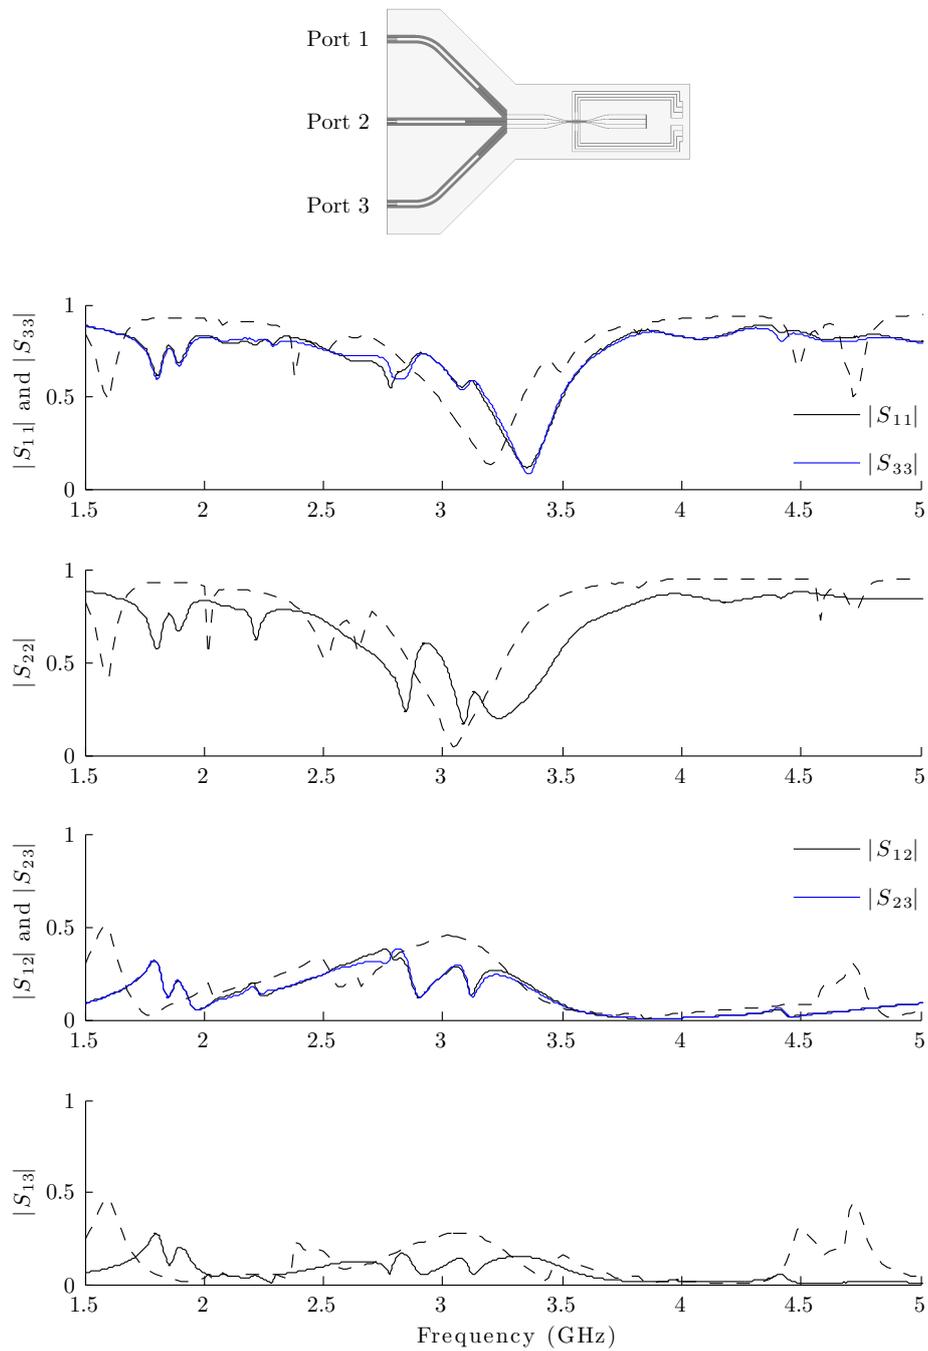


Figure 8.8: Comparison of VNA data (solid line) and the HFSS simulation (dashed line) of the total fraction of power incident on a given microwave port that is coupled into the trap. The VNA data has had effects of the measured in-vacuum cabling and feedthrough losses subtracted to enable better comparison with the model.

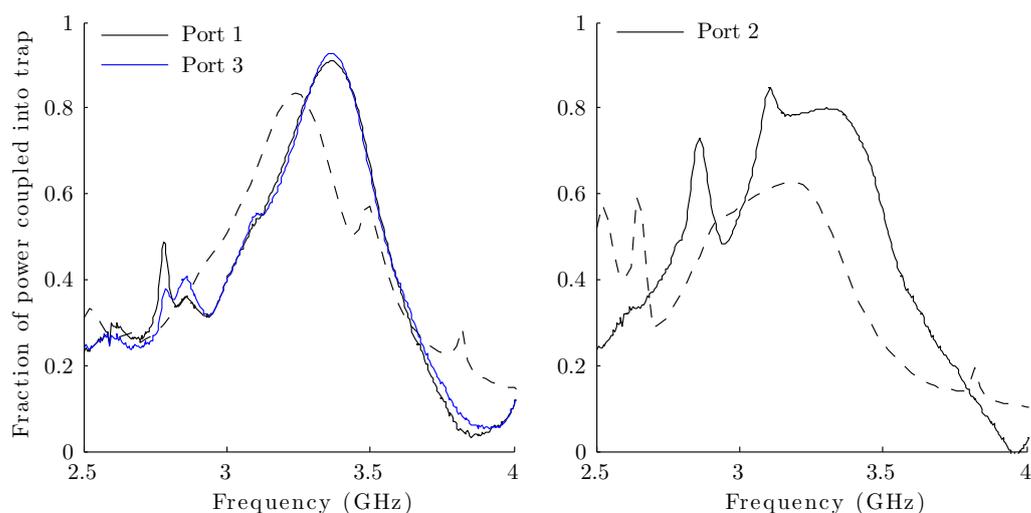


Figure 8.9: Comparison of VNA data (solid line) and the HFSS simulation (dashed line). The VNA data has had the independently measured losses of the in-vacuum cabling and feedthrough subtracted to allow for better comparison with the model.

8.2.1.1 Fabrication tolerances

Because the trap dimensions are photolithographically defined we expect the fractional error on large dimensions like the lengths of the microwave resonators to be negligible. The gaps between electrodes close to the ion are only $10\ \mu\text{m}$ though. Because of the polarity of our mask (see section 2.4) any scattering of light under the mask or erosion of the resist will decrease this width. From optical microscopy we estimate the gap to be $> 8\ \mu\text{m}$ and we re-simulate the trap using this worst-case value (see figure 8.10). The other dimension that is relatively poorly defined is the electrode thickness which varies by at most $\pm 1\ \mu\text{m}$. This is due to variations in solution flow and electric field lines across the trap during electroplating. We also re-simulated the trap with varying thickness in figure 8.10. Neither effect appears large enough to explain the discrepancies between the model and data.

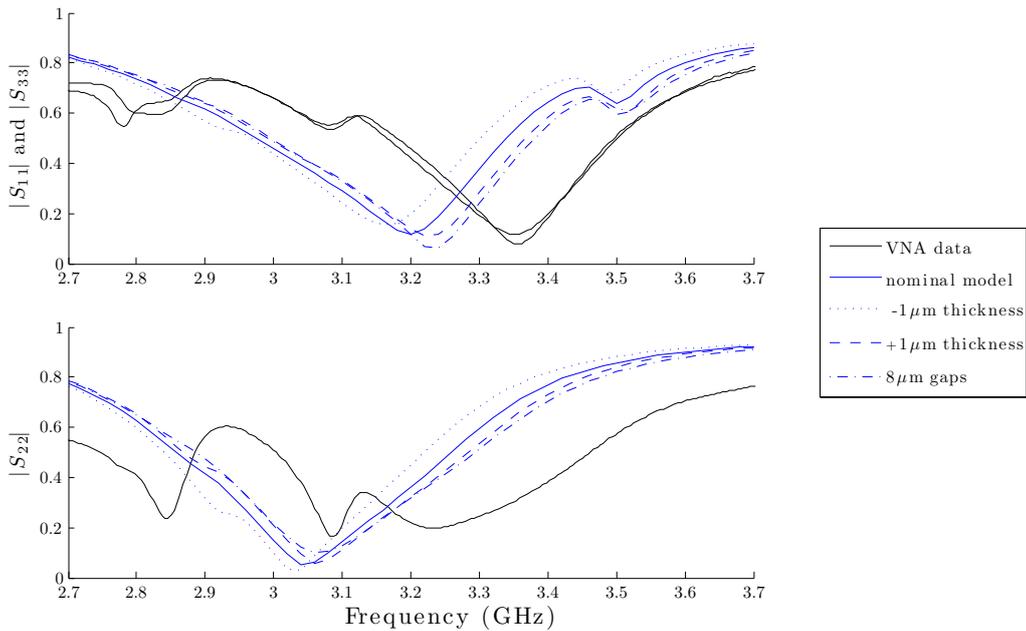


Figure 8.10: Modelled effect of varying the electrode thickness (nominally $5\ \mu\text{m}$) and the inter-electrode gap width (nominally $10\ \mu\text{m}$) to simulate possible errors in the trap fabrication. The VNA data and nominal model are also shown.

8.2.2 RF/microwave diplexer

Figure 8.11 shows the finished diplexer. The two diplexer boards are mounted in separate brass boxes bolted back-to-back (see figure 8.11). The boards are glued down to the base of the box with conductive epoxy to ensure mechanical stability and a good thermal contact. The SMA connectors¹⁷ are panel-mounted with an extended PTFE dielectric which extends through the box wall. Figure 8.12 shows the comparison between the Microwave Office (MWO) model and VNA data from the assembled diplexer circuit. The level of agreement is qualitatively very good, however the return loss never drops below $\sim 10\ \text{dB}$. The probable explanation for this is reflections from the coaxial to microstrip transition, which are not included in the model, rather than a problem with the diplexer circuit. These reflections are undesirable because they set up resonances between the diplexer and the trap. We are currently redesigning the diplexer box with better

¹⁷Amphenol RF, 901-9891-RFX

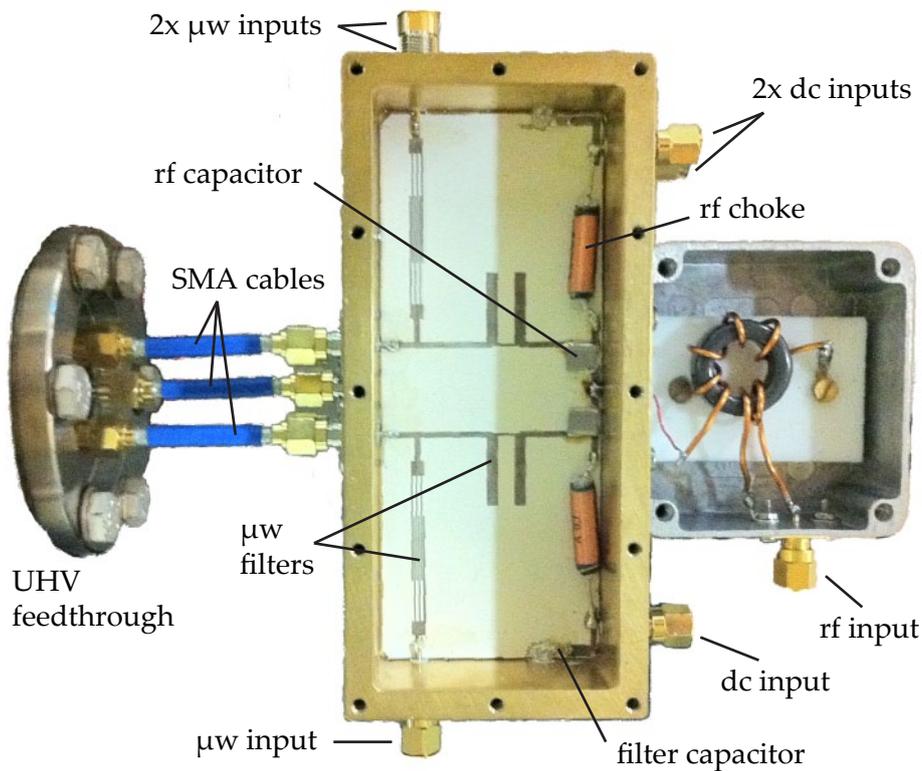


Figure 8.11: Photograph of assembled diplexer with metal lids removed. The third diplexer channel for the centre trap electrode is inverted with respect to the two channels shown and is identical except that the rf connection is grounded.

optimisation of these transitions.

8.2.3 Drive system

8.2.3.1 IQ mixer calibration

In order to simplify nulling of the field we wish to calibrate the IQ mixers such that we can program the DACs in terms of phase and amplitude rather than I and Q which are non-linear (see figure 8.13). We do this by recording the phase and amplitude at various values of I and Q to create a lookup table. The amplitude is measured using a microwave power meter¹⁸ and the phase is measured by using a mixer as a phase detector.

¹⁸Agilent V3500A, ± 0.21 dB absolute accuracy

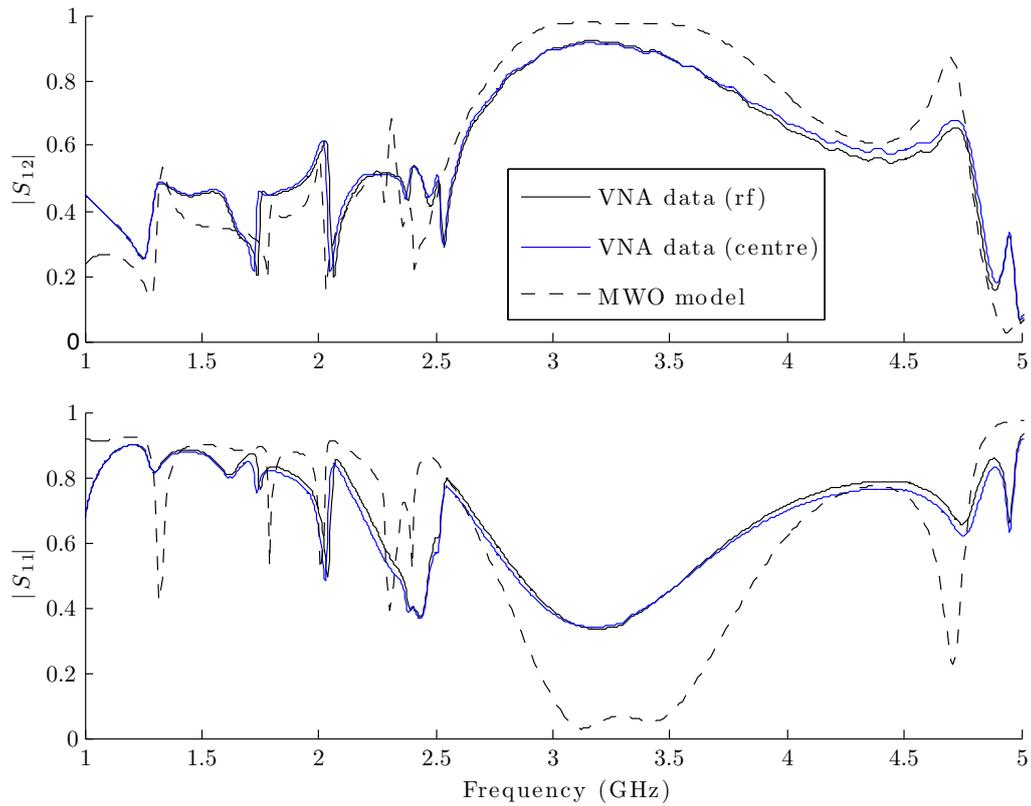


Figure 8.12: Comparison of MWO circuit simulation of the rf/microwave diplexer with measured VNA data. VNA data is shown for one of the diplexer channels which is connected to the rf transformer (rf) and for the channel diplexer with no rf connection (centre).

8.2.3.2 Microwave phase and amplitude stability

The extent to which we can null the qubit carrier transition will be limited by the quality and positional stability of our magnetic field null at the ion and the positional stability of the ion (see section 8.3.3). The quality of the null is a function of the microwave phase stability as a null only exists if the three microwave currents are exactly in phase with each other. The positional stability of the null is a function of the relative amplitude of the three currents (see figure 7.2). We can measure the extent to which the phase and amplitude can be controlled by inserting a 90° hybrid (see figure 8.14) at the point in the drive system where the diplexer would be, with one input connected to each of two microwave channels (see figure 8.15). We terminate one output port and measure how well we

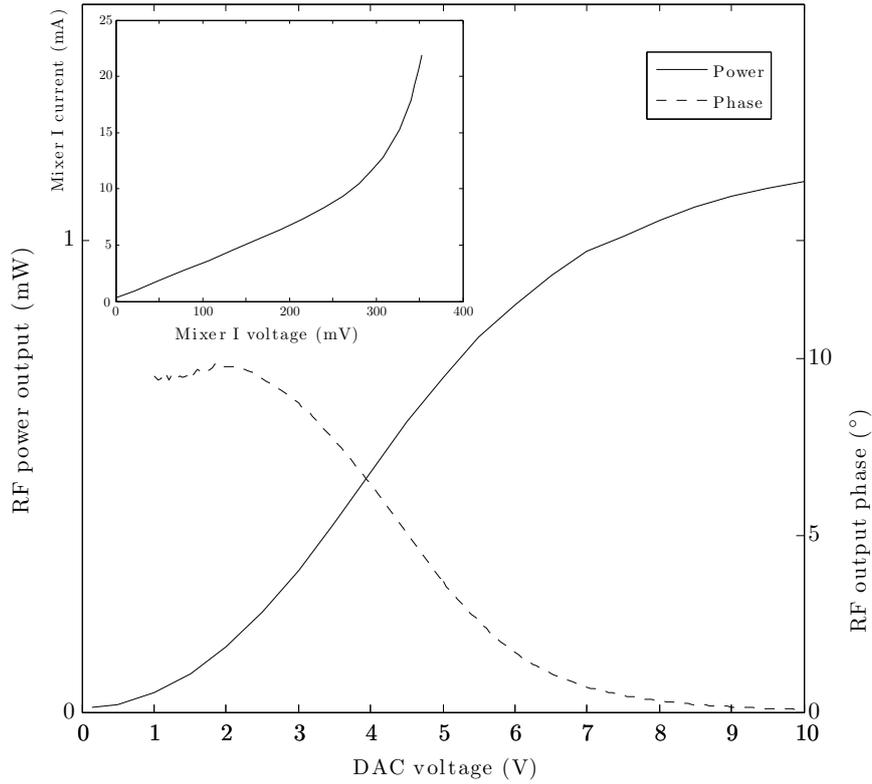


Figure 8.13: Power output and phase as a function of DAC voltage for one input of an IQ mixer (with $P_{LO} = 10$ dBm). The inset shows the current-voltage curve for the mixer.

can null the output of other output port by adjusting the microwave phase and amplitude of one drive channel using the IQ mixer. The ratio of the two output powers is given by

$$\frac{P_2}{P_3} = \frac{|V_1^+ + jV_4^+|^2}{|jV_1^+ + V_4^+|^2}. \quad (8.2)$$

We cannot differentiate between phase and amplitude error but if we assume no phase error (taking phase error into account does not change the final result) then the ratio of input powers can be calculated as

$$\frac{P_1}{P_4} = \left(\frac{1 - \sqrt{\frac{P_2}{P_3}}}{1 + \sqrt{\frac{P_2}{P_3}}} \right)^2. \quad (8.3)$$

On the trap itself we will be combining three channels rather than two but this measurement will still give us a good indication of the drive system performance. In addition this performance will be measured into a good 50Ω load but

8. CONSTRUCTION AND TESTING OF MICROWAVE-DRIVEN QUANTUM LOGIC TRAP

$$\begin{array}{l}
 V_1^+ \text{ --- } \boxed{\text{X}} \text{ --- } V_2^- = \frac{1}{\sqrt{2}}(V_1^+ + jV_4^+) \\
 V_4^+ \text{ --- } \boxed{\text{X}} \text{ --- } V_3^- = \frac{1}{\sqrt{2}}(jV_1^+ + V_4^+)
 \end{array}$$

Figure 8.14: The output voltage amplitudes (V_2^- and V_3^-) as a function of input voltage amplitudes (V_1^+ and V_4^+) for an ideal 90° hybrid.

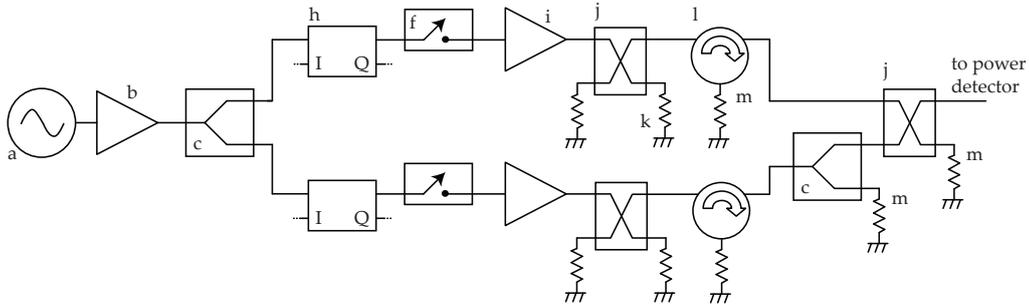


Figure 8.15: Subsection of the sideband microwaves drive (see figure 7.7) with the two outputs feeding a 90° hybrid. Parts are as in table 7.3.

we expect similar performance into a slightly mismatched load such as the trap because the last stage of the drive system is an isolator which diverts reflected power away from the drive system.

Measurements were first made without pulsing the microwaves (CW mode) using the microwave power meter (with a low-pass filter in to block harmonics) on port 2 of the hybrid. In CW we were able adjust one of the IQ mixers to null the carrier to a level where there was 57 dB less power on the port 2 than port 3. We also used a spectrum analyser to look at the nulled signal as this is a very sensitive way of detecting noise away from the carrier. The only resolved peaks were sidebands at 100 Hz (presumably from the 50 Hz mains) at ~ -50 dB below the 3.2 GHz carrier leakage and sidebands at 70 kHz (source unknown) at ~ -30 dB below, neither of which are anticipated to be an issue.

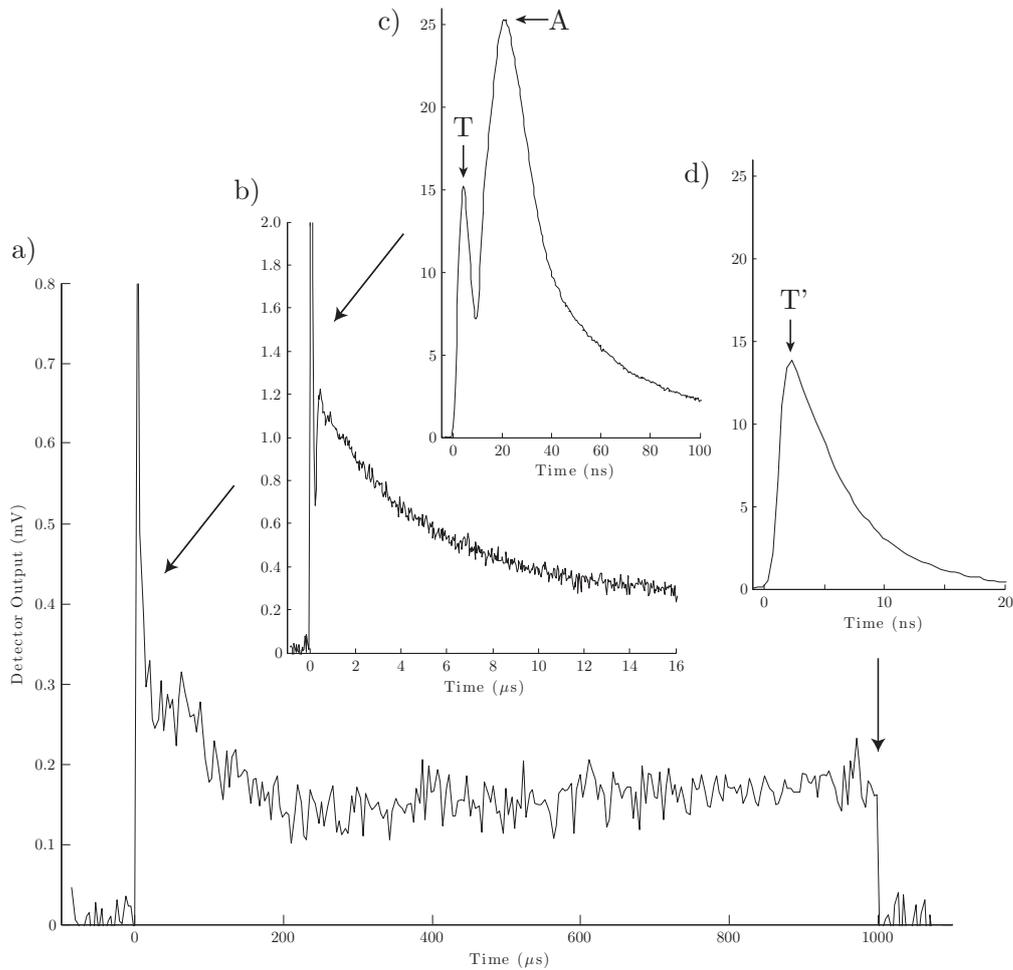


Figure 8.16: Graph a) shows the output of port 2 of the hybrid over a 1 ms pulse. Graphs b) and c) show zoomed in views of the first 16 μs and 100 ns of the pulse respectively. Graph d) shows the last 20 ns of the pulse (the zero of graph d is shifted by 1 ms). Note the different vertical axis ranges. The transients T and T' are from imperfect temporal overlap and the transient A is from amplifier non-linearity. The detector has a non-linear response but was calibrated at relevant points using a power meter.

8.2.3.3 Pulsed phase and amplitude stability

To measure the phase and amplitude stability during pulses (1 ms pulses at 50 Hz repetition rate) we replaced the power meter by a diode detector¹⁹ connected to a high speed oscilloscope²⁰. Figure 8.16 shows that it is not possible to null to the CW level due to transients during the pulse. We believe these are due to thermal effects and the inductance of current supplies in the power amplifiers. Note that one channel was set to a power level of 1.6 W and the other 3.2 W (with half the power diverted away from hybrid using a power splitter). This was done to ensure we are not benefitting from common-mode suppression of thermal transients that only occurs when operating both amplifiers at exactly the same power. The nulling is better than 30 dB during the first 50 ns, 40 dB during the first 10 μ s and 50 dB for the rest of the pulse (1 ms). After half an hour of running the nulling over the bulk of the pulse was still better than 49 dB.

8.2.3.4 Pulse temporal overlapping

As well as matching the phase and amplitude, one also has to equalise the time delay through the two channels of the drive system to ensure that the pulses are temporally overlapped. This is done to first order by using equal lengths of cabling and identical components throughout each channel. It can then be fine-tuned by adjusting the relative lengths of the cabling to the TTL inputs of the microwave switches. We were unable to perform this nulling perfectly, presumably due to slight variations in rise and fall times (typical 10%-90% time of 5 ns) of the switches causing slight pulse length differences. This residual temporal mismatch is the cause of the first peak in figure 2.8c (marked T) and the peak at the end of the pulse in figure 2.8d (marked T').

¹⁹Agilent 423B, 8-12 ns rise time into 50 Ω

²⁰Tektronix TDS 5104, 1 GHz, 5 GS/s

8.2.3.5 Projected gate error

We assume we there are two microwave fields of opposite phase at the ion, $B_{\mu w,a}$ and $B_{\mu w,b}$ with input powers $P_a \propto B_{\mu w,a}^2$ and $P_b \propto B_{\mu w,b}^2$ respectively. The residual $B_{\mu w}$ -field is given by

$$B_{\mu w} = B_{\mu w,a} - B_{\mu w,b} = B_{\mu w,a} \left(1 - \frac{B_{\mu w,b}}{B_{\mu w,a}} \right) = B_{\mu w,a} \left(1 - \sqrt{\frac{P_b}{P_a}} \right) \quad (8.4)$$

where $\frac{P_b}{P_a}$ is calculated using equation 8.3. Noting that the Rabi frequency Ω_R is proportional to $B_{\mu w}$ we can use equations 6.14 and 6.15 to calculate both the ac Zeeman shift and the off-resonant excitation we expect from $B_{\mu w}$. The ac Zeeman shift is dominated by the 50 dB nulling error over the bulk of pulse which implies $P_b/P_a = 0.987$. For the anticipated microwave field amplitude for the gate, $B_{\mu w,a} = 500 \mu\text{T}$, this gives a residual $B_{\mu w} = 3.2 \mu\text{T}$ and a Rabi frequency due to the residual field of $\Omega_R = 2\pi \times 44 \text{ kHz}$. This in turn gives an ac Zeeman shift of $\Omega_R^2/4\Delta = 2\pi \times 120 \text{ Hz}$ at $\Delta = 2\pi \times 4 \text{ MHz}$. The off-resonant excitation is dominated by the transient at the start of the pulse (marked A in figure 2.8c) during which the nulling is 30 dB, implying $P_b/P_a = 0.88$. This gives a residual $B_{\mu w} = 31 \mu\text{T}$ and a Rabi frequency due to the residual field of $\Omega_R = 2\pi \times 430 \text{ kHz}$. For $t_{pulse} = 20 \text{ ns}$ and $\Delta = 2\pi \times 4 \text{ MHz}$ the off resonant excitations is 7.1×10^{-4} . These errors are small enough that demonstrating entanglement should be still possible but show that further work on developing the drive system is required.

As both of these effects scale with Rabi frequency squared (in this regime) only modest technical improvements are required to make these effects negligible. The transients could be reduced passively, or by using local feedback at the amplifier, or by using pulse shaping. Passive improvement could be achieved by using more powerful class A amplifiers at well below saturation. This would help because in that regime the output transistor operates linearly around a bias point and so does not draw more current or dissipate more power during the pulse. This approach, whilst simple, is expensive and inefficient.

Feedback could be implemented by wrapping the power amplifiers in feedback loops which fix the gain and phase shift introduced by the amplifier. Making this feedback loop fast enough and stable would be a complex task but it is documented in the literature [Lee03]. Finally, if AWGs were used to drive the IQ mixers (see section 7.2.1.1) we could reduce the fast transients by turning the pulses on more smoothly over a longer period. Then we could characterise the slower transients and feed forward on the pulse envelope shape to remove them. We propose to implement the final scheme first because the AWG control will give us added flexibility to implement more complex pulse shaping.

Another possible factor that may limit how well the microwave field at the ion can be nulled is the resonator on the trap itself. This is because during the pulse losses in the resonator will cause it to heat up, increasing its resistance and decreasing the circulating power, and this effect will be different for the three resonators. In principle these transient effects could be modelled in the HFSS simulation but we decided not to invest the time necessary to do that as there is little that can be done to mitigate it at this stage. Instead we will investigate this by observing the effect on the ion directly.

8.2.3.6 Output power

The maximum power output from each channel of the drive system is 35 dBm (3.2 W). At this point the power amplifier output level is ~ 39 dBm, which is further than 3 dB below its 1 dB compression point.

8.3 Testing with trapped ions

8.3.1 Ion loading and lifetime

Single ions and two-ion crystals of either $^{40}\text{Ca}^+$ or $^{43}\text{Ca}^+$ are loaded easily at trap secular frequencies of 4 MHz radial and 500 kHz axial. Single ion life-

time is around one hour and the two-ion lifetime appears to be consistent with 30 mins, though this has not been carefully measured. Background scatter is ~ 200 counts/s at with $2.0 \mu\text{W}$ of 397 nm power (in a spot with $1/e^2$ radius $w = 25 \mu\text{m}$), compared to of order 50,000 counts/s from a cold ion close to resonance.

8.3.2 Motional frequency stability

The stability of the motional modes in the trap is an important parameter as any drifts will cause an asymmetrical error in the gate detuning (see section 6.6.3.2). In order to characterise this drift we measured the frequency of one of the radial modes of a single ion over 110 minutes (see figure 8.17) using the method described below. The drift over a few tens of minutes is around 1 kHz; any drifts over longer timescales will be straightforward to measure and correct for. This 1 kHz drift is unacceptably high as it is comparable to our gate detuning δ . Whilst measuring the drift we also measured the rf power going into the transformer using a directional coupler. The fluctuations on this power can be converted to a fractional frequency fluctuation since the radial frequencies are approximately $\propto \sqrt{P_{rf}}$ and when plotted on the same graph as the mode frequency appear to explain the drift. The temperature in the lab was also strongly correlated with the fluctuations. Thus by better temperature stabilising the rf circuitry or implementing a feedback loop on the rf power we hope to reduce frequency fluctuations to well below 1 kHz.

8.3.2.1 Improved ‘tickle’ procedure

The secular motion of an ion is often found by applying an rf voltage to one of the dc electrodes and sweeping its frequency across those of the ion’s secular modes. On resonance this voltage heats the ion and, if the ion is being Doppler cooled, causes a change in the average fluorescence rate. A disadvantage of this

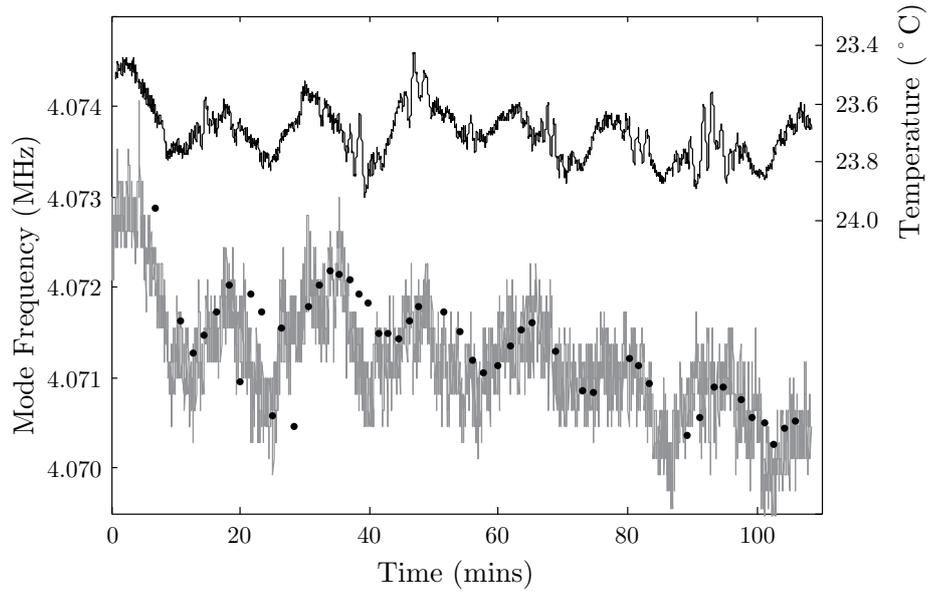


Figure 8.17: Points show the measured frequency of one of the radial modes of a single ion over 110 minutes (circles). The ion was lost and a new ion loaded at $t=10$ minutes. The grey graph overlaid is the expected frequency from an independent measure of the power going into the rf transformer. The upper black line is the temperature (note that the scale is inverted), as measured by a sensor in the rf power detector.

method is that the amount of heating that is required for it to be visible is quite large. This broadens out the resonance and causes hysteresis effects where the measured mode frequency depends on the direction in which the frequency is swept. There is also a concern that in a surface trap the anharmonic pseudopotential explored by large orbits may shift the frequencies. Experimentally we find that this technique is limited to a resolution of about 1 kHz.

A more sensitive technique is to use the photon correlation equipment used for micromotion compensation (see sections 2.2.1.1 and 3.4) to correlate photon arrival time with the phase of the applied rf voltage. This is done by simply taking the TAC trigger from this synthesiser rather than the trap rf drive synthesiser. The phase of the photon arrival histogram at several points through the resonance can then be fitted. Experimentally we find the rf voltages required are about an order of magnitude smaller and the resolution is increased to about 100 Hz (in comparable measurement time).

8.3.3 Micromotion

The connections to the two rf electrodes on the trap are made via separate halves of the diplexer board and separate feedthroughs and cabling (see section 7.2.5). We were concerned that this could lead to a phase shift between the two voltages by the time they reached the trap which would cause uncompensatable micromotion [BMB⁺98]. Every effort was made to keep both rf paths as identical as possible. Experimentally the uncompensatable micromotion we see is close to the limit of our detection sensitivity (~ 2 nm amplitude) which is negligible.

The micromotion compensation indicates that the E-field offset at the ion is stable to better than ± 5 V/m over a day (with no detectable change when loading an ion), which corresponds to a displacement of ± 18 nm. A displacement of 18 nm from the microwave null with a 10 T/m gradient gives $\mathbf{B}_{\mu w} = 0.18 \mu\text{T}$. The ac Zeeman shift from this field at a detuning of 4 MHz is < 1 Hz.

8.4 Microwave field measurements

Initial measurements of the microwave field were made by observing Rabi flops of a single $^{43}\text{Ca}^+$ ion at $B_0 = 23$ G. After a brief experimental optimisation of cooling laser parameters we obtained a PMT signal of 20000 counts/s from the ion. This is $\sim 25\%$ of the maximum obtainable from $^{40}\text{Ca}^+$ and is comparable with previous work on $^{43}\text{Ca}^+$ at lower fields (see section 6.3). A more careful optimisation of parameters may however yield higher fluorescence.

8.4.1 Microwave field measurements by Rabi flopping

By measuring the Rabi frequencies of a π and a σ microwave transition we are able to calculate the magnitude and direction of $\mathbf{B}_{\mu w}$ when applying microwaves to each trap electrode in turn (if we assume $B_{\mu w}^y = 0$). The measurement procedure begins with initialisation to the $M_m = +4$ state using the

397 nm σ beam set to σ^+ polarisation. A microwave pulse resonant with the $S_{1/2}, M_m = +4$ to $M_n = +3$ σ^- transition was then applied for a varying length of time (transition a in figure 8.18). After this pulse $\sim 90\%$ of the population in $M_m = +4$ is shelved to the $D_{5/2}$ level using the 393 nm σ^+ beam. Finally the state is detected with the cooling lasers; a bright ion indicating projection to the $M_n = +3$ state. Fitting the population as a function of microwave pulse length to equation 6.15 gives the Rabi frequency (see figure 8.19).

In order to measure the Rabi frequency of the π transition b the experiment is repeated but with a Rabi π pulse on transition a immediately preceding the variable length pulse on transition b . The shelving procedure remains unchanged as the 393 nm laser will shelve the $M_m = +3$ state almost as well as the $M_m = +4$. The microwave power input into the diplexer was 1 mW which, when the loss of the diplexer, feedthrough and cabling is taken into account, corresponds to 0.747 mW at the trap. Table 8.2 shows the measured $B_{\mu w}$ -fields and the simulated values for that power input. The $B_{\mu w}$ -field was measured at 3.169 and 3.177 GHz while the simulation is carried out at 3.200 GHz (the frequency for which it was optimised). The comparison with theory is somewhat complicated by the fact the actual resonance frequency of the trap is not the same as the simulation; however the fact that we are already achieving better than the simulated values whilst slightly off-resonance bodes well for achieving the designed field gradient.

Table 8.2: Measured B-fields at 3.169 and 3.177 GHz with 1 mW input to the diplexer (0.747 mW at the trap) compared with the values from the HFSS simulation (at 3.200 GHz). The simulation assumes \mathbf{B}_0 is parallel to the trap.

Port	Measured $\mathbf{B}_{\mu w}$ (μT)	Angle to \mathbf{B}_0 ($^\circ$)	Simulated $\mathbf{B}_{\mu w}$ (μT)	Angle to \mathbf{B}_0 ($^\circ$)
1	16.9	77.4	14.5	64.7
2	19.0	12.6	17.3	0
3	14.7	69.2	14.5	64.7

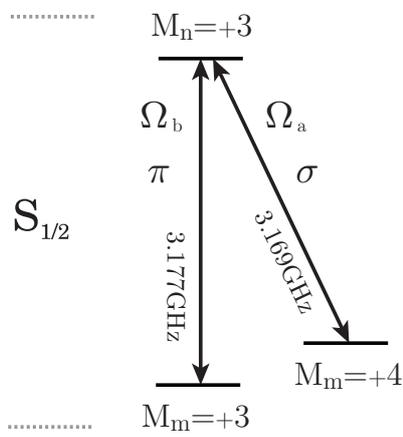


Figure 8.18: The two microwave transitions which we used to measure Rabi frequencies at a static field $B_0 = 23 \text{ G}$. The $M_m = +4$ to $M_n = +3$ σ^- transition is driven by $B_{\mu w}^z$; the $M_m = +3$ to $M_n = +3$ π transition is driven by $B_{\mu w}^x$.

8. CONSTRUCTION AND TESTING OF MICROWAVE-DRIVEN QUANTUM LOGIC TRAP

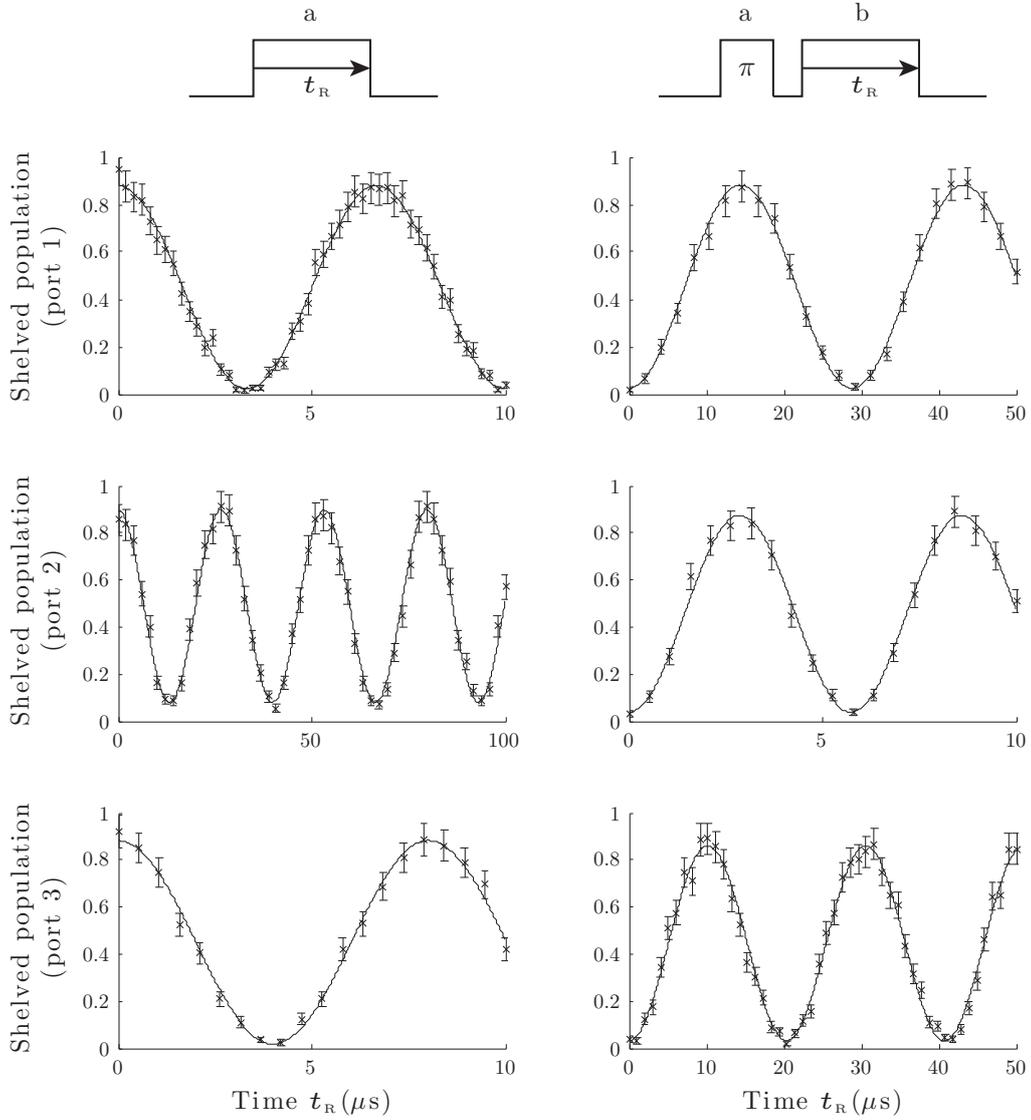


Figure 8.19: Rabi flops on two transitions when microwave power is applied to each electrode in turn. Fitted Rabi frequencies are $\Omega_{a,1} = 2\pi \times 150$ kHz, $\Omega_{b,1} = 2\pi \times 34.7$ kHz, $\Omega_{a,2} = 2\pi \times 37.6$ kHz, $\Omega_{b,2} = 2\pi \times 174$ kHz, $\Omega_{a,3} = 2\pi \times 125$ kHz and $\Omega_{b,3} = 2\pi \times 49.1$ kHz.

9

Conclusions

I was the first guy playing Daft Punk to the rock kids.
I played it at CBGB's.
Everyone thought I was crazy.
We all know.
I was there.

LCD Soundsystem

9.1 Summary

In chapter 3 we described a new electrode design for a surface-electrode Paul trap, which allows rotation of the normal modes out of the trap plane, and a technique for micromotion compensation in all directions using a two-photon process, which avoids the need for an ultraviolet laser directed to the trap plane. The fabrication and characterisation of the trap were described, as well as its implementation for the trapping and cooling of single Ca^+ ions. We also proposed a repumping scheme that increases ion fluorescence and simplifies heating rate measurements obtained by time-resolved ion fluorescence during Doppler cooling.

In chapter 4 we characterised the performance of a surface-electrode trap fabricated by Sandia National Laboratories using established semiconductor integrated circuit and micro-electro-mechanical-system (MEMS) microfabrication

processes, which are in principle scalable to much larger ion trap arrays, as proposed for implementing ion trap quantum information processing. We measured rf ion micromotion parallel and perpendicular to the plane of the trap electrodes, and find that on-package capacitors reduce this to <10 nm in amplitude. We also measured ion-trapping lifetime, charging effects due to laser light incident on the trap electrodes, and the heating rate for a single trapped ion. The performance of this trap is found to be comparable with others of the same size scale.

In chapter 5 laser-cleaning of the electrodes of the trap used in chapter 4 was attempted using ns pulses from a tripled Nd:YAG laser at 355 nm. The effect of the laser pulses at several energy-density levels was tested by measuring the heating-rate of an ion as a function of its secular frequency. A reduction of the electric-field noise spectral density by $\sim 50\%$ was observed and a change in the frequency dependence also noticed. This is the first reported experiment where the anomalous heating phenomenon has been reduced by removing the source as opposed to reducing its thermal driving by cryogenic cooling. This technique may open the way to better control of the electrode surface quality in ion microtraps.

In chapter 6 we gave an overview of hyperfine structure in the ground state of $^{43}\text{Ca}^+$ and identified three potential qubits suitable for a microwave-driven multi-qubit logic gate. One qubit is a low field qubit that has previously been implemented and the other two are intermediate-field qubits. We described both the advantages and some of the technical issues that come with using an ion with low-lying D-levels in the intermediate field. Techniques for qubit initialization, readout, single and multi-qubit gates were described.

In chapter 7 the required specifications of a trap that incorporates microwave circuitry to drive quantum logic gates were outlined. Simulations were used to design an electrode layout that best meets these specifications, as well as a

half-wave resonant structure and associated coupling elements that reduce the microwave power required to implement quantum logic. A microwave drive system capable of delivering the microwave pulses for driving the gate was also designed.

Chapter 8 described construction of the trap and experimental apparatus designed in chapter 7. Testing of sections of the microwave drive system indicate that it will be capable of driving gates at close to the required fidelity, though further development is still needed. Experiments to characterise the trap show that one- and two-ion $^{43}\text{Ca}^+$ crystals can be loaded and cooled and that their motional heating rate is low. Ion initialisation and readout have been demonstrated and single-qubit microwave gates show that the microwave field at the ion is in close agreement with simulations. None of the characterisations performed so far seems to pose a major limitation for high fidelity microwave-driven entanglement.

9.2 Towards microwave-driven entanglement

The immediate goal in the group is to continue the work in chapter 8 to the point where microwave-driven entanglement can be demonstrated. The following points describe the experimental steps that need to be taken.

- Implement the full qubit initialisation drive to reach the field-independent qubit transition.
- Align the 850 nm σ beam onto the ion and optimise initialisation and readout of the qubit.
- Use a Ramsey experiment to measure qubit decoherence time.
- Connect one sideband drive to the trap in order to apply microwave pulses to all three electrodes at once. Attempt to null $\mathbf{B}_{\mu w}$ at the ion. Measure the

gradient by recording the ac Zeeman shift as a function of ion position as demonstrated in [OWC⁺11].

- Stabilise the drift in the radial trapping frequency.
- Demonstrate single sideband Rabi flops.
- By probing the relative red and blue sideband strengths measure the initial temperature and heating rates of one and two ions.
- Implement the other sideband drive system and demonstrate two-qubit entanglement.
- Investigate fidelity limitations. If limited by the control and drive systems then make improvements to them to improve the fidelity.

9.2.0.1 High-field qubits

We envisage using the 23G qubit initially, but we will work concurrently on demonstrating the high-field qubits. At the point that this is successful we will be able to switch between the qubits as we choose; switching to the other qubit transitions is trivial as it only requires adjustment of the microwave frequencies.

9.3 Future development: the 50-qubit processor

In the rest of the conclusion I will describe how the Sandia ion trap architecture described in chapters 4-5 could be combined with the $^{43}\text{Ca}^+$ hyperfine qubit and the microwave-driven multi-qubit gates in chapters 6-8 to create a quantum information processor. This processor would be able to carry out operations with errors below the fault-tolerant threshold on a non-trivial number of qubits (~ 50). For each part of the processor I will outline whether the necessary experimental techniques have already been demonstrated or require further development.

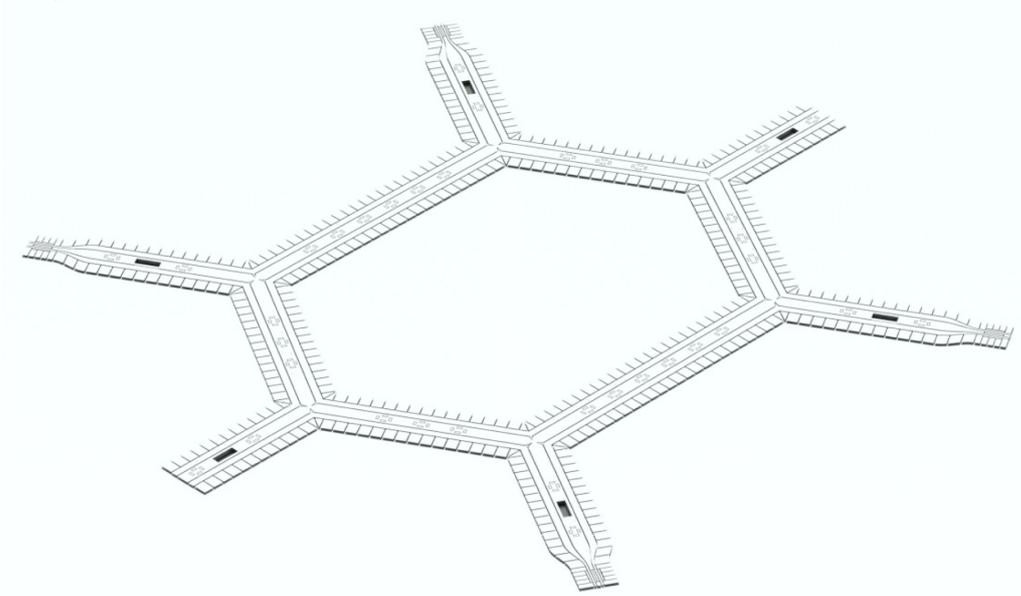


Figure 9.1: The hypothetical 50-qubit processor.

9.3.1 Trap

The basic operation of the processor is as outlined by Kielpinski, Monroe and Wineland [KMW02]. In this scheme information is moved around the processor by physically shuttling the ions around using control electrodes.

The design of the processor (see figure 9.1) is based on the surface trap array experimentally demonstrated by Amini and coworkers at NIST Boulder [AUW⁺10]. Improved Y-junctions (10^6 shuttles through junction without loss) as well as ion separation and recombination have since been demonstrated using the Sandia architecture [MHS⁺11].

9.3.1.1 Cryogenic operation

Cryogenic cooling of the processor will be crucial for several reasons.

Much recent work with ion traps has concentrated on liquid helium (4 K) cooling; however liquid nitrogen cooling appears to be favourable for this type of processor. Primarily this is because we require Watts of cooling power due to the microwave power dissipated. Due to its much lower latent heat of va-

porisation (2.6 kJ/l vs. 160 kJ/l) only milliWatts of cooling power is available from a He cryostat. Superconductors offer a possible solution to this issue and would allow operation at He temperatures, though they are not entirely lossless at microwave frequencies and it is not clear whether they can support the required current densities (this could however be simulated). Superconducting resonators are currently under development for experiments involving trapped ionic molecules above surface traps [McC10].

At room temperature the lowest achievable pressures are typically around 10^{-12} Torr and will limit lifetime to of order a few hours in a typical surface trap. Whilst this is acceptable for traps working with a few ions, this loss rate would be inconvenient for a 50 ion processor. Collisions also reorder multiple ion crystals which is problematic [HHJ⁺09] and cryogenic cooling will reduce this significantly.

At 77K the resistance of the aluminium electrodes will decrease by about an order of magnitude, allowing a factor of ~ 3 increase in current for the same power dissipation. From figure 9.2 we see that the thermal conductivity of the silicon substrate also increases by an order of magnitude (note that silicon is already a factor of 3 better than sapphire at room temperature). Together these two effects will allow for a factor of 20 increase in current (and thus gate speed) for the same temperature rise. Unfortunately the thermal advantages may not be realised in the current Sandia architecture as the SiO_2 dielectric layer has a relatively poor thermal conductivity. Replacing the lowest SiO_2 layer with one of a higher thermal conductivity may however be possible [BCM⁺05].

Cryogenic cooling will reduce the heating rate. From [Lab08] we might expect a Temperature dependence of $S_E(\propto 1 + (T/T_0)^\beta$ where $T_0 \approx 10 - 70$ K and $\beta \approx 2 - 4$. Thus much of the heating rate reduction has already been obtained by the time the trap has been cooled to 77 K.

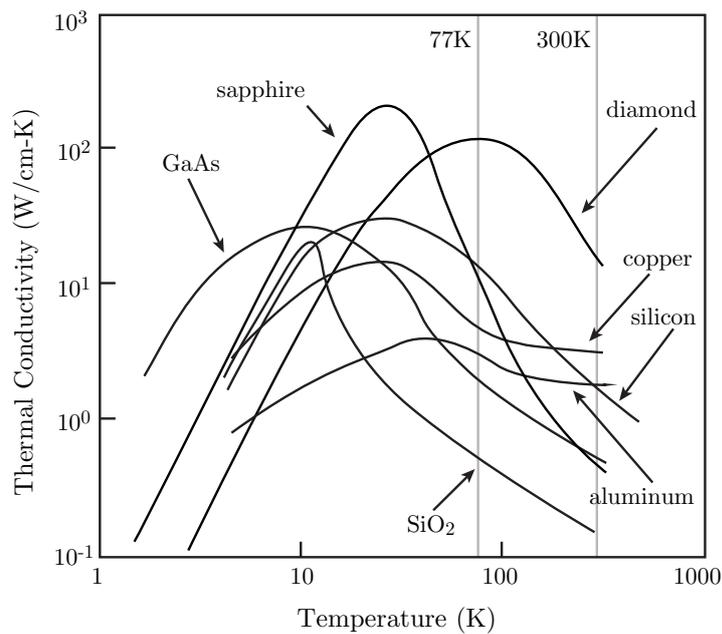


Figure 9.2: Thermal conductivities of several materials used in microfabrication (figure from [GDDC01]).

9.3.2 Laser interaction region

Integration of optics into the trap will be important for scaling as it will remove the need for optical access to the cryostat and alignment of poorly scalable bulk optics. It will also eliminate photon scatter off parts of the trap other than the laser interaction regime, reducing trap charging.

Integration of fibre optics and diffractive optics into the ion trap for fluorescence detection has been demonstrated in a Sandia trap nominally identical to the one in chapter 4 [BEM⁺11]. Laser beam delivery through the same techniques should also be possible in the future and has already been demonstrated in a simpler setup without focussing optics [KHC11]. These laser beams will be used to initialise and read out the ion as described in chapter 6.

Only microwatt level laser beam powers are required at each laser interaction region so one set of diode lasers would be sufficient to run the processor envisaged.

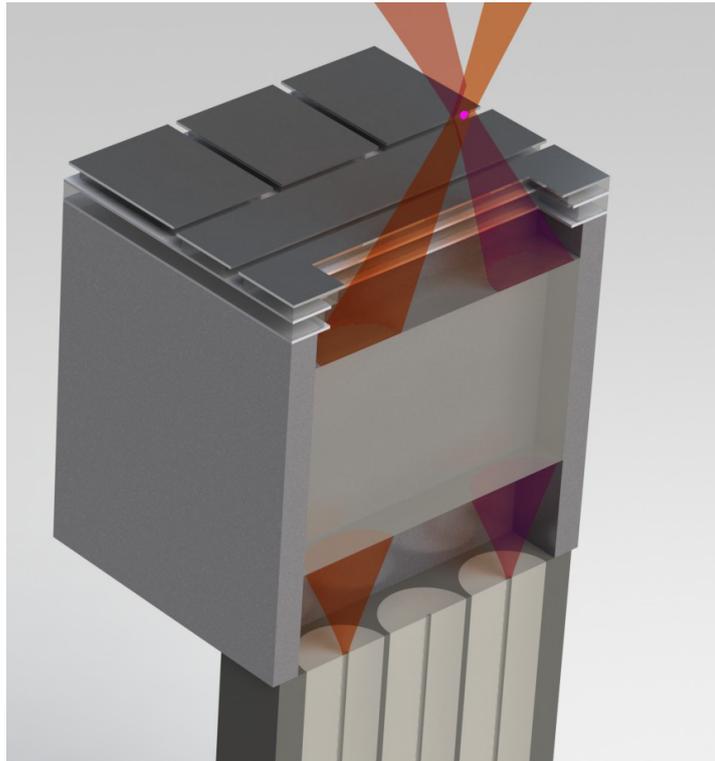


Figure 9.3: Cross-section of a laser interaction zone. The fibre on the left delivers IR light, the fibre on the right delivers UV light and the centre fibre collects fluorescence from the ion. The central insert is the fused silica optics die which has the diffractive optics and polarisers fabricated onto it.

9.3.2.1 Sympathetic cooling

In order to carry out the multi-qubit gates the ions need to be laser cooled close to the Doppler limit. It is likely that this will need to be done before each gate operation as any ions that have been stored or moved for some time are likely to have heated by several quanta.

Before initialisation and after readout the qubit ions can be laser cooled directly. However once they are in use as a qubit, any scattered photons will destroy the qubit. A solution to this is to use a different species of ion to sympathetically cool the qubit ions [KKM⁺00].

There are two possible solutions. The first is to use another isotope of calcium and rely on the isotope shift to avoid decoherence of the qubit isotope via

photon scattering (as shown in figure 4.6). Unfortunately the isotope shift is relatively small ($^{42}\text{Ca}^+$ is the optimal isotope with ~ 1.2 GHz shift from the closest $^{43}\text{Ca}^+$ transition) so for cooling on the dipole transition quite tight focussing of the 397 nm beam on the coolant ion would be required to achieve low enough scattering [LRH⁺04, HMS⁺09]. The quadrupole transitions at either 729 nm or 733 nm with a much narrower linewidth could be used to reduce this error but at the expense of increased laser complexity and power requirements and reduced cooling rate [HSC⁺08]. The other option would be to use a completely different element, allowing for simple cooling with no errors. For $^{43}\text{Ca}^+$ the best choice of sympathetic cooling species appears to be $^{88}\text{Sr}^+$. This is because out of the ions compatible with integrated optics it has the most similar mass to $^{43}\text{Ca}^+$ (see table 6.1) which is important for efficient cooling [KKM⁺00]. Another advantage is that as the cooling transition is 422 nm compared to 397 nm for Ca^+ it is possible the same set of integrated optics could be used for both species.

Shuttling of a two-ion crystal through a junction [Bla10] and ion crystal splitting [BCS⁺04] have both been achieved experimentally with < 1 quanta of heating. Thus it should be possible to laser-cool the ion crystal in the laser interaction region, split it from the refrigerant ion (though this step is not a requirement [HHJ⁺09]) and then shuttle it to the gate region.

9.3.3 Single-qubit microwave operations

Figure 9.5 shows how a set of microwave electrodes could be integrated into the trap structure. If every ion in the processor were stored above such a zone then single-qubit operations could be carried out on every qubit in the processor in parallel. These operations have already been demonstrated in a single-zone trap with errors of $2.0(2) \times 10^{-5}$ [BWC⁺11] (well below the fault-tolerant threshold) with no fundamental limit on further improvement. For multi-zone traps crosstalk is an additional issue and this is discussed below. The ability to carry

out single qubit operations in parallel with negligible error would allow the use of dynamical decoupling techniques to reduce qubit decoherence. This will require of order 200 microwave drive channels but, as the power requirements are modest (~ 1 mW), each channel would consist of only a small number of integrated circuits on a circuit board and it would be straightforward to scale to this number of channels. This should be contrasted to laser operations with Raman laser gates which have yet to be demonstrated for errors below 4.8×10^{-3} [KLR⁺08]. Whilst decreasing this error is possible in principle [OIB⁺07], the power and beam pointing stabilities required are challenging. Beyond that scaling to 50 ions would be non-trivial as we would need to deliver a pair of laser beams at mW power levels to every trapping location.

9.3.4 Crosstalk

The one major disadvantage of microwave control is that when microwaves are applied to an electrode on the processor they will produce some finite microwave field across the entire processor, leading to erroneous operations on qubits other than the target qubit. There are three ways this crosstalk could be reduced to acceptable levels. Firstly the crosstalk could be reduced through optimisation of the microwave design. By using the Sandia architecture all the microwave feed lines could be led in as striplines underneath the upper ground plane. They would emerge through this ground plane only directly underneath the target ion (see figure 9.5). The size and shape of the electrodes would be optimised to produce the largest possible field at the ion in that region relative to the other regions in the trap. Secondly the field induced by each microwave electrode at all the other trapping locations should be entirely reproducible and could be carefully measured. This would allow the fields at these locations to be accounted for in the pulses applied to these ions, effectively nulling them [OLA⁺08]. Finally composite pulse sequences could be used to protect against

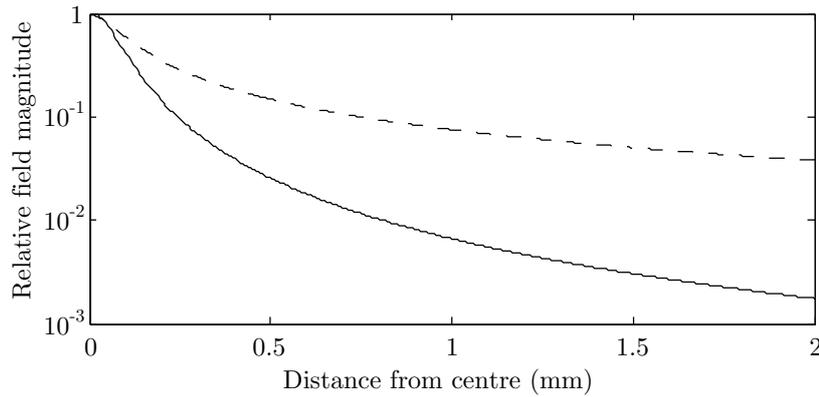


Figure 9.4: The microwave field as a function of distance d from a coplanar wave guide transmission line at a constant height above the substrate of $75\ \mu\text{m}$ (solid line). The magnitude of the field is relative to its value directly above the CPW. The CPW dimensions are $30\ \mu\text{m}$ centre conductor width and $10\ \mu\text{m}$ gap either side. The substrate is $500\ \mu\text{m}$ sapphire with $5\ \mu\text{m}$ gold metallisation. The dashed line shows how the field magnitude would look if it were proportional to $1/d$.

crosstalk errors [HRB08], even if the nulling pulses are subject to noise.

9.3.4.1 Crosstalk simulation

To estimate what level of crosstalk might be expected between two different gate regions on the trap we simulated the microwave field as a function of distance d from a coplanar wave guide transmission line (see figure 9.4). As the microwave field drops off rapidly (much more rapidly than $1/d$) it appears that for single-qubit gates in traps spaced a few hundred μm apart nulling pulses will only have to be accurate to the 10^{-2} level in order to reduce crosstalk to below 10^{-4} . Nulling to this level should be technically straightforward. The microwave magnetic fields used during the multi-qubit gates are somewhat larger than for single qubit gates; however they are detuned from the qubit frequency. To give an illustrative example, the field during the gate 1 mm from the gate region in the trap of chapter 7 is $\sim 10\ \mu\text{T}$. This would give an ac Zeeman shift (see section 6.2.4) of $\sim 4\ \text{kHz}$. If this field could again be nulled to the 10^{-2} level this ac Zeeman shift would be reduced to $\sim 0.4\ \text{Hz}$.

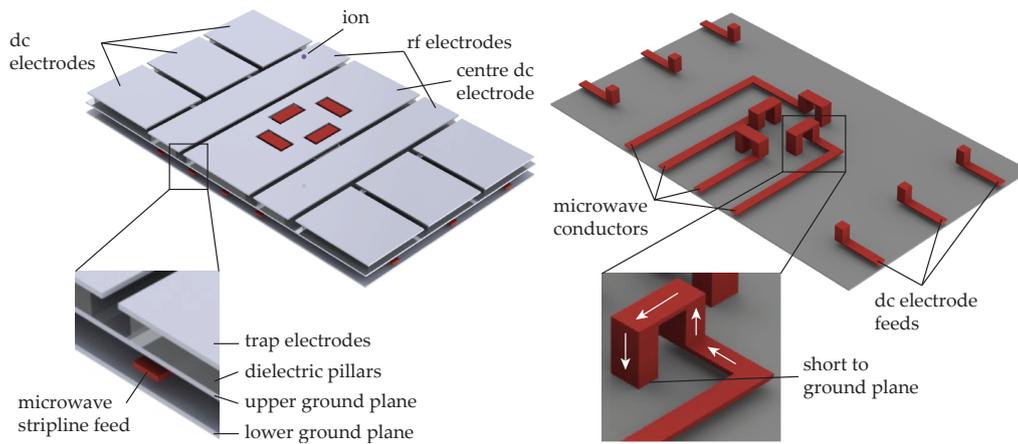


Figure 9.5: Trap zone for single qubit rotations and qubit storage (left) and the same zone with the trap electrodes and upper groundplane removed (right). The lower left inset shows the flow of the microwave currents (white arrows). The lower right inset shows how the two ground planes and the microwave conductor form a stripline waveguide.

9.3.5 Multiple qubit microwave operations

During the gate there will be some crosstalk with ions not in the gate region. By storing ions not involved in the gates in the single-qubit zones the same techniques of crosstalk reduction as used in single-qubit gates could also be applied here. There is also the added advantage that the gate drive frequencies are detuned from the carrier by the secular frequency.

9.3.5.1 Increasing the gate speed

As well as using cryogenic cooling to increase the current density possible, a significant increase in gate speed will require a much smaller ion-electrode distance d and electrode dimensions. A factor of 4 (from $75 \mu\text{m}$ to $\sim 19 \mu\text{m}$) decrease in d will increase the gate speed by a factor of 16 for the same current ($B'_{\mu w} \propto d^{-2}$). As there will also be a factor 4 decrease in electrode width the current density will increase (but by less than factor of 4 due to current crowding at the edge of the electrodes). Along with the current increases that cryogenic operation will allow, $10 \mu\text{s}$ gate times appear feasible.

At this ion-electrode distance it is likely that ion heating will be much more

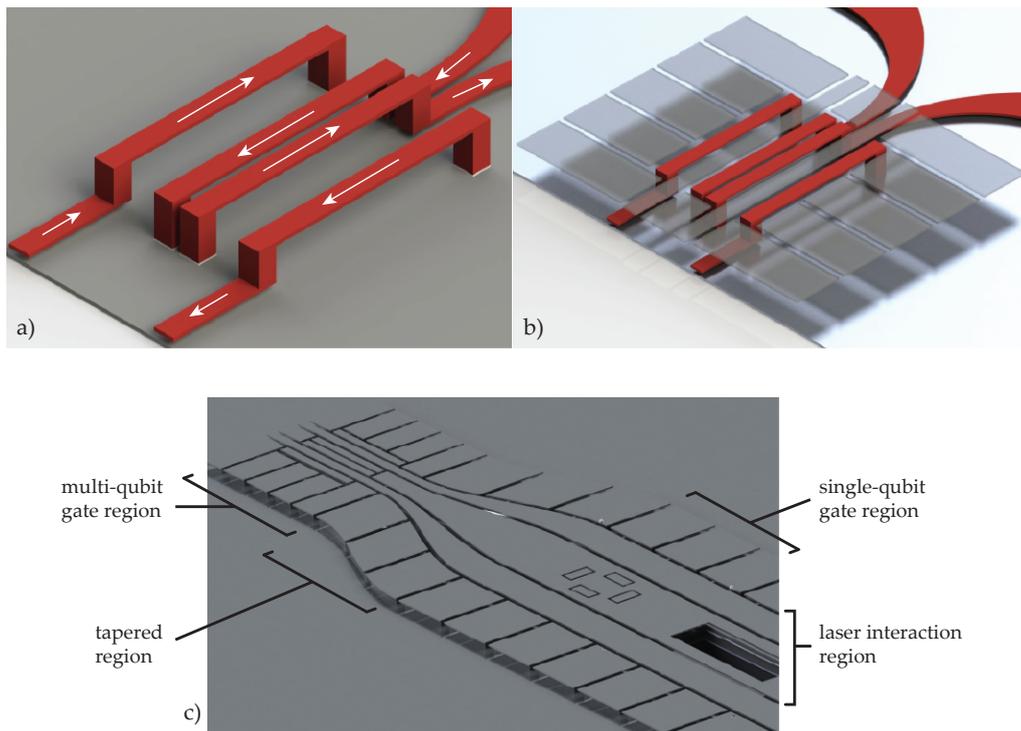


Figure 9.6: The multiple-qubit gate region of the trap. a) The microwave current-carrying electrodes (red) with arrows to show direction of current flow. One end of each electrode is grounded to the lower ground plane grey. b) The same electrodes but also showing the trap electrodes (transparent) The upper ground plane is not shown. c) The gate region shown connected to the rest of the processor via a tapered region.

of a problem. Cryogenic cooling will certainly mitigate the problem somewhat, though to what extent is still not well understood. Hopefully future development of the laser cleaning techniques introduced in chapter 5 or further understanding of the origin of the electric field noise will allow heating rates to be reliably lowered such that that gates can be carried out in traps of this size.

9.3.5.2 Resonators

The multi-level Sandia architecture opens up more possibilities with regards to the microwave layout than the single layer we were constrained to in chapter 7. One thing that is now possible is the incorporation of vias to the ground plane which would allow for smaller, capacitively-coupled, quarter-wave resonators. Figure 9.7a shows how two such resonators could be incorporated

with the microwave conductor layout 4 in section 7.1.3.1 (only two resonators are used as the third microwave conductor carries relatively little current). By meandering the resonators, there would be enough space for all the resonators on a modestly-sized processor. A more detailed design study would be required to ascertain the best solution.

9.3.6 Scaling further

Whilst 50 qubits appears feasible scaling to much larger arrays would be difficult due to the density of microwave signals and control voltages required. This may be alleviated by integrating much of the control electronics onto the processor itself. Another promising approach would be to connect the processor to a photonics quantum network. This would allow for scalability by combining many processors. A summary of how remote ion qubits can be entangled via photons can be found in [DM10].

9.3.7 Conclusion

In order to realise the processor described most of the required techniques have been demonstrated or will require only modest technical improvements. However five main areas still need significant development.

- Increase the fidelity of the multi-qubit microwave gates beyond the fault-tolerant threshold. By further development of the work in chapters 7-8 we hope to address this.
- Demonstrate single-qubit addressing with near-field microwaves. We have started designing a proof-of-concept experiment to demonstrate this.
- Reduction of anomalous heating rate, possibly by using laser cleaning techniques similar to those described in this thesis.

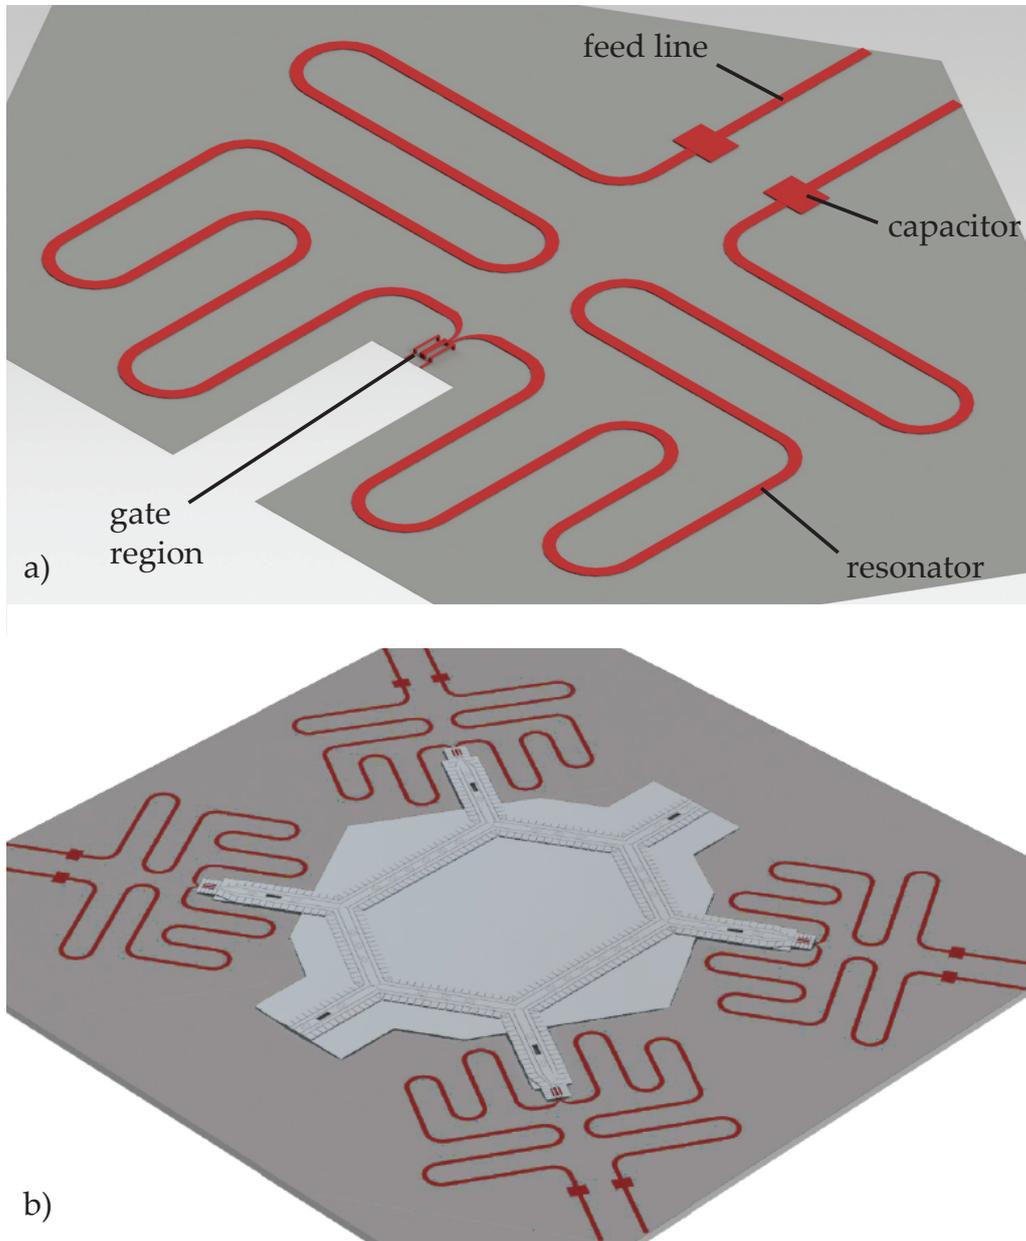


Figure 9.7: a) A zoomed out image of the microwave electrodes in figure 9.6a showing the resonator structures and coupling capacitors. b) Illustration of how the resonators for the four gate zones could be incorporated onto the processor chip. The light grey upper ground plane is partially suppressed to show the resonator layout.

9. CONCLUSIONS

- Combination of microwave elements with a multilayer trap fabrication technology.
- Demonstrate integrated optics that can both deliver laser light of the correct polarization and collect fluorescence from the ion.

Given the intensive research efforts worldwide on developing trapped ion quantum computing, such a processor appears technically feasible in the 5-10 year timescale.

Bibliography

- [ABG⁺94] F. Arbes, M. Benzing, Th. Gudjons, F. Kurth, and G. Werth. Precise determination of the ground state hyperfine structure splitting of ^{43}Ca II. *Zeitschrift für Physik D: Atoms, Molecules and Clusters*, 31:27, 1994. [cited at p. 91]
- [ABLW11] J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. *Atom Chips*, chapter Microfabricated chip traps for ions, pages 395–420. Wiley-VCH, 2011. [cited at p. 7, 40, 73]
- [AHJ⁺11] D. T. C. Allcock, T. P. Harty, H. A. Janacek, N. M. Linke, C. J. Ballance, A. M. Steane, D. M. Lucas, R. L. Jarecki Jr., S. D. Habermehl, M. G. Blain, D. Stick, and D. L. Moehring. Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap. *arXiv*, 1105.4864v1 [quant-ph], 2011. [cited at p. 73, 84]
- [All07] D. T. C. Allcock. A multi-channel laser diagnostic system for trapped-ion quantum computing experiments. Master’s thesis, University of Oxford, 2007. [cited at p. 20]
- [AUW⁺10] J. M. Amini, H. Uys, J. H. Wesenberg, S. Seidelin, J. Britton, J. J. Bollinger, D. Leibfried, C. Ospelkaus, A. P. VanDevender, and D. J. Wineland. Toward scalable ion traps for quantum information processing. *New Journal of Physics*, 12:033031, 2010. [cited at p. 10, 173]
- [BB02] D. J. Berkeland and M. G. Boshier. Destabilization of dark states and optical spectroscopy in Zeeman-degenerate atomic systems. *Physical Review A*, 65:033413, 2002. [cited at p. 94, 97]
- [BBHK00] M. G. Boshier, G. P. Barwood, G. Huang, and H. A. Klein. Polarisation-dependent optical pumping for interrogation of a magnetic-field-independent “clock” transition in laser-cooled trapped $^{87}\text{Sr}^+$. *Applied Physics B: Lasers and Optics*, 71:51, 2000. [cited at p. 29]
- [BCL⁺07] K. R. Brown, R. J. Clark, J. Labaziewicz, P. Richerme, D. R. Leibbrandt, and I. L. Chuang. Loading and characterization of a printed-circuit-board atomic ion trap. *Physical Review A*, 75:015401, 2007. [cited at p. 39, 56]
- [BCM⁺05] N. Bresson, S. Cristoloveanu, C. Mazuré, F. Letertre, and H. Iwai. Integration of buried insulators with high thermal conductivity in SOI MOS-FETs: Thermal properties and short channel effects. *Solid-State Electronics*, 49(9):1522, 2005. [cited at p. 174]
- [BCS⁺04] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Deterministic quantum teleportation of atomic qubits. *Nature*, 429:737, 2004. [cited at p. 8, 177]
- [BEM⁺11] G. R. Brady, A. R. Ellis, D. L. Moehring, D. Stick, C. Highstrete, K. M. Fortier, M. G. Blain, R. A. Haltli, A. A. Cruz-Cabrera, R. D. Briggs, J. R.

- Wendt, T. R. Carter, S. Samora, and S. A. Kemme. Integration of fluorescence collection optics with a microfabricated surface electrode ion trap. *Applied Physics B: Lasers and Optics*, 103(4):801, 2011. [cited at p. 40, 61, 76, 78, 117, 175]
- [BGG⁺03] G. P. Barwood, K. Gao, P. Gill, G. Huang, and H. A. Klein. Observation of the hyperfine structure of the $2S_{1/2}$ - $2D_{5/2}$ transition in $^{87}\text{Sr}^+$. *Physical Review A*, 67:013402, 2003. [cited at p. 90]
- [BHL⁺05] K.-A. Brickman, P. C. Haljan, P. J. Lee, M. Acton, L. Deslauriers, and C. Monroe. Implementation of Grover's quantum search algorithm in a scalable system. *Physical Review A*, 72:050306, 2005. [cited at p. 6]
- [Bia11] R. C. Bialczak. *Development of the fundamental components of a superconducting qubit quantum computer*. PhD thesis, University of California Santa Barbara, 2011. [cited at p. 133]
- [BKRB08a] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Experimental quantum-information processing with $^{43}\text{Ca}^+$ ions. *Physical Review A*, 77:062306, 2008. [cited at p. 107]
- [BKRB08b] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, 4(6):463, 2008. [cited at p. 5, 112]
- [Bla10] R. B. Blakestad. *Transport of trapped-ion qubits within a scalable quantum processor*. PhD thesis, University of Colorado, 2010. [cited at p. 71, 177]
- [BMB⁺98] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83:5025, 1998. [cited at p. 43, 49, 54, 67, 165]
- [BMC⁺98] D. J. Berkeland, J. D. Miller, F. C. Cruz, B. C. Young, J. C. Bergquist, W. M. Itano, and D. J. Wineland. High-accuracy microwave and optical frequency standards using laser-cooled Hg^+ ions. *Laser Physics*, 8:673, 1998. [cited at p. 90]
- [BTLT09] E. A. Burt, S. Taghavi-Larigani, and R. L. Tjoelker. High-resolution spectroscopy of $^{201}\text{Hg}^+$ hyperfine structure: A sensitive probe of nuclear structure and the hyperfine anomaly. *Physical Review A*, 79:062506, 2009. [cited at p. 90]
- [Bur10] A. H. Burrell. *High fidelity readout of trapped ion qubits*. PhD thesis, University of Oxford, 2010. [cited at p. 47, 107]
- [BWC⁺11] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Physical Review A*, 84:030303, 2011. [cited at p. 4, 105, 177]
- [CBB⁺05] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Surface-electrode architecture for ion-trap quantum information processing. *Quantum Information and Computation*, 5:419, 2005. [cited at p. 39, 124]
- [CBL⁺05] J. Chiaverini, J. Britton, D. Leibfried, E. Knill, M. D. Barrett, R. B. Blakestad, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, T. Schaetz, and D. J. Wineland. Implementation of the semiclassical quantum Fourier transform in a scaleable system. *Science*, 308:997, 2005. [cited at p. 6]

- [CL08] J. Chiaverini and W. E. Lybarger. Laserless trapped-ion quantum simulations without spontaneous scattering using microtrap arrays. *Physical Review A*, 77:022324, 2008. [cited at p. 113]
- [CLS⁺04] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland. Realization of quantum error correction. *Nature*, 432:602, 2004. [cited at p. 6]
- [CMQ⁺10] W. C. Campbell, J. Mizrahi, Q. Quraishi, C. Senko, D. Hayes, D. Hucul, D. N. Matsukevich, P. Maunz, and C. Monroe. Ultrafast gates for single atomic qubits. *Physical Review Letters*, 105:090502, 2010. [cited at p. 4]
- [Cur10] M. J. Curtis. *Measurement-selected ensembles in trapped ion qubits*. PhD thesis, University of Oxford, 2010. [cited at p. 16, 73]
- [CZ95] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20):4091, 1995. [cited at p. 2]
- [DCG09] R. Dubessy, T. Coudreau, and L. Guidoni. Electric field noise above surfaces: A model for heating-rate scaling law in ion traps. *Physical Review A*, 80:031402, 2009. [cited at p. 76, 84]
- [DHK⁺83] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B: Lasers and Optics*, 31(2):97, 1983. [cited at p. 15]
- [DHL⁺04] L. Deslauriers, P. C. Haljan, P. J. Lee, K-A. Brickman, B. B. Blinov, M. J. Madsen, and C. Monroe. Zero-point cooling and low heating of trapped ¹¹¹Cd⁺ ions. *Physical Review A*, 70:043408, 2004. [cited at p. 75]
- [DiV05] D. P. DiVincenzo. *The Physical Implementation of Quantum Computation*, pages 1–13. Wiley-VCH Verlag GmbH, 2005. [cited at p. 1]
- [DKN⁺10] M. R. Dietrich, N. Kurz, T. Noel, G. Shu, and B. B. Blinov. Hyperfine and optical barium ion qubits. *arXiv*, 1004.1161v1 [quant-ph], 2010. [cited at p. 90]
- [DM10] L.-M. Duan and C Monroe. Colloquium: Quantum networks with trapped ions. *Reviews of Modern Physics*, 82:1209, 2010. [cited at p. 182]
- [DNM⁺11] N. Daniilidis, S. Narayanan, S. A. Möller, R. Clark, T. E. Lee, P. J. Leek, A. Wallraff, S. Schulz, F. Schmidt-Kaler, and H. Häffner. Fabrication and heating rate study of microscopic surface electrode ion traps. *New Journal of Physics*, 13:013032, 2011. [cited at p. 58, 71, 76, 80, 84]
- [DO06] P. Delaporte and R. Oltra. Laser cleaning: state of the art. In J. Perriere, E. Millon, and E. Fogarassy, editors, *Recent advances in laser processing of materials*, European Materials Research Society Series, pages 411–440. Elsevier Amsterdam, 2006. [cited at p. 77]
- [DOS⁺06] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. Scaling and suppression of anomalous heating in ion traps. *Physical Review Letters*, 97:103007, 2006. [cited at p. 71, 75]
- [DRS⁺10] L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Preparation and measurement of three-qubit entanglement in a superconducting circuit. *Nature*, 467:574, 2010. [cited at p. 2]

- [ESL⁺07] R. J. Epstein, S. Seidelin, D. Leibfried, J. H. Wesenberg, J. J. Bollinger, J. M. Amini, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Simplified motional heating rate measurements of trapped ions. *Physical Review A*, 76:033411, 2007. [cited at p. 71, 75]
- [FSG⁺08] A. Friedenauer, H. Schmitz, J. T. Glueckert, D. Porras, and T. Schaetz. Simulating a quantum magnet with trapped ions. *Nature Physics*, 4:757, 2008. [cited at p. 90]
- [GDDC01] E. A. Gutiérrez-D., M. J. Deen, and C. L. Claeys. *Low temperature electronics: physics, devices, circuits, and applications*. Academic Press, 2001. [cited at p. 175]
- [GKH95] V. Gusev, A. A. Kolomenskii, and P. Hess. Effect of melting on the excitation of surface acoustic wave pulses by UV nanosecond laser pulses in silicon. *Applied Physics B: Lasers and Optics*, 61:285, 1995. [cited at p. 79]
- [GKW⁺04] S. Groth, P. Krüger, S. Wildermuth, R. Folman, T. Fernholz, D. Mahalu, I. Bar-Joseph, and J. Schmiedmayer. Atom chips: fabrication and thermal properties. *Applied Physics Letters*, 85(14):2980, 2004. [cited at p. 120, 126]
- [GRL⁺03] S. Gulde, M. Riebe, G. P. T. Lancaster, C. Becher, J. Eschner, H. Häffner, F. Schmidt-Kaler, I. L. Chuang, and R. Blatt. Implementing the Deutsch-Jozsa algorithm on an ion-trap quantum computer. *Nature*, 422:408, 2003. [cited at p. 5, 6]
- [HBHB10] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. Trapped-ion probing of light-induced charging effects on dielectrics. *New Journal of Physics*, 12:093035, 2010. [cited at p. 69, 82]
- [HCD⁺11] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, Q. Quraishi, and C. Monroe. Coherent error suppression in spin-dependent force quantum gates. *arXiv*, 1104.1347v1 [quant:ph], 2011. [cited at p. 113]
- [HHJ⁺09] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. Complete methods set for scalable ion trap quantum information processing. *Science*, 325:1227, 2009. [cited at p. 5, 174, 177]
- [HHJ⁺11] J. P. Home, D. Hanneke, J. D. Jost, D. Leibfried, and D. J. Wineland. Normal modes of trapped ions in the presence of anharmonic potentials. *New Journal of Physics*, 13:073026, 2011. [cited at p. 9]
- [HLB⁺05] P. C. Haljan, P. J. Lee, K-A. Brickman, M. Acton, L. Deslauriers, and C. Monroe. Entanglement of trapped-ion clock states. *Physical Review A*, 72:062316, 2005. [cited at p. 90]
- [HMM⁺10] D. Hayes, D. N. Matsukevich, P. Maunz, D. Hucul, Q. Quraishi, S. Olmschenk, W. Campbell, J. Mizrahi, C. Senko, and C. Monroe. Entanglement of atomic qubits using an optical frequency comb. *Physical Review Letters*, 104:140501, 2010. [cited at p. 4, 5]
- [HMS⁺09] J. P. Home, M. J. McDonnell, D. J. Szwer, B. C. Keitch, D. M. Lucas, D. N. Stacey, and A. M. Steane. Memory coherence of a sympathetically cooled trapped-ion qubit. *Physical Review A*, 79:050305, 2009. [cited at p. 66, 177]
- [Hou08] M. G. House. Analytic model for electrostatic fields in surface-electrode ion traps. *Physical Review A*, 78:033402, 2008. [cited at p. 117]

- [HRB08] H. Häffner, C. F. Roos, and R. Blatt. Quantum computing with trapped ions. *Physics Reports*, 469:155, 2008. [cited at p. 3, 5, 179]
- [HRW07] D. B. Hume, T. Rosenband, and D. J. Wineland. High-fidelity adaptive qubit detection through repetitive quantum nondemolition measurements. *Physical Review Letters*, 99:120502, 2007. [cited at p. 4]
- [HSC⁺08] R. J. Hendricks, J. L. Sørensen, C. Champenois, M. Knoop, and M. Drewsen. Doppler cooling of calcium ions using a dipole-forbidden transition. *Physical Review A*, 77:021401, 2008. [cited at p. 177]
- [IEK⁺11] R. Islam, E. E. Edwards, K. Kim, S. Korenblit, C. Noh, H. Carmichael, G.-D. Lin, L.-M. Duan, C.-C. Joseph Wang, J. K. Freericks, and C. Monroe. Onset of a quantum phase transition with a trapped ion quantum simulator. *Nature Communications*, 2:377, 2011. [cited at p. 6]
- [ITU11] Y. Ibaraki, U. Tanaka, and S. Urabe. Detection of parametric resonance of trapped ions for micromotion compensation. *Applied Physics B: Lasers and Optics*, 2011. [cited at p. 56]
- [Jam98] D. F. V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B: Lasers and Optics*, 66:181, 1998. [cited at p. 111]
- [KBZ⁺09a] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and R. Blatt. Deterministic entanglement of ions in thermal states of motion. *New Journal of Physics*, 11:023002, 2009. [cited at p. 112]
- [KBZ⁺09b] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and R. Blatt. High-fidelity entanglement of $^{43}\text{Ca}^+$ hyperfine clock states. *Physical Review A*, 79:020304, 2009. [cited at p. 94]
- [Kei07] B. C. Keitch. *A quantum memory qubit in calcium-43*. PhD thesis, University of Oxford, 2007. [cited at p. 17, 24, 104]
- [KHC11] T. H. Kim, P. F. Herskind, and I. L. Chuang. Surface-electrode ion trap with integrated light source. *Applied Physics Letters*, 98:214103, 2011. [cited at p. 175]
- [KKM⁺00] D. Kielpinski, B. E. King, C. J. Myatt, C. A. Sackett, Q. A. Turchette, W. M. Itano, C. Monroe, D. J. Wineland, and W. H. Zurek. Sympathetic cooling of trapped ions for quantum logic. *Physical Review A*, 61:032310, 2000. [cited at p. 176, 177]
- [KLR⁺08] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Josta, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77:012307, 2008. [cited at p. 4, 178]
- [KMK⁺04] E. Koukharenko, Z. Moktadir, M. Kraft, M. E. Abdelsalam, D. M. Bagnall, C. Vale, M. P. A. Jones, and E. A. Hinds. Microfabrication of gold wires for atom guides. *Sensors and Actuators A*, 115(2-3):600, 2004. [cited at p. 31, 32]
- [KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709, 2002. [cited at p. 39, 173]

- [KWM⁺99] B. E. King, C. S. Wood, C. J. Myatt, Q. A. Turchette, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland. Cooling the collective motion of trapped ions to initialize a quantum register. *Physical Review Letters*, 81(7):1525, 1999. [cited at p. 112]
- [Lab08] J. Labaziewicz. *High fidelity quantum gates with ions in cryogenic microfabricated ion traps*. PhD thesis, Massachusetts Institute of Technology, 2008. [cited at p. 31, 42, 174]
- [Lan06] C. E. Langer. *High Fidelity Quantum Information Processing with Trapped Ions*. PhD thesis, University of Colorado, 2006. [cited at p. 3]
- [LAS⁺11] N. M. Linke, D. T. C. Allcock, D. J. Szwer, C. Ballance, T. Harty, H. Janacek, D. N. Stacey, A. M. Steane, and D. M. Lucas. Background-free detection of trapped ions. *Applied Physics B: Lasers and Optics*, 2011. [cited at p. 66]
- [Lee03] J. L. Lee. *Feedback Linearization of RF Power Amplifiers*. PhD thesis, Stanford University, 2003. [cited at p. 162]
- [Lev86] M. H. Levitt. Composite pulses. *Progress in NMR Spectroscopy*, 18:61, 1986. [cited at p. 104]
- [LGA⁺08] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang. Suppression of heating rates in cryogenic surface-electrode ion traps. *Physical Review Letters*, 100:013001, 2008. [cited at p. 40, 71]
- [LGL⁺08] J. Labaziewicz, Y. Ge, D. R. Leibbrandt, S. X. Wang, R. Shewmon, and I. L. Chuang. Temperature dependence of electric field noise above gold surfaces. *Physical Review Letters*, 101:180602, 2008. [cited at p. 75, 76]
- [LHC11] G. H. Low, P. F. Herskind, and I. L. Chuang. Finite geometry models of electric field noise from patch potentials in ion traps. *arXiv*, 1109.2995v1 [quant-ph], 2011. [cited at p. 76, 84]
- [LHN⁺11] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos. Universal digital quantum simulation with trapped ions. *Science*, 334:57, 2011. [cited at p. 6]
- [LJL⁺10] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien. Quantum computers. *Nature*, 464:45, 2010. [cited at p. 2]
- [LKC⁺05] C. Lisowski, M. Knoop, C. Champenois, G. Hagel, M. Vedel, and F. Vedel. Dark resonances as a probe for the motional state of a single ion. *Applied Physics B: Lasers and Optics*, 81(1):5, 2005. [cited at p. 53]
- [LKH⁺07] D. M. Lucas, B. C. Keitch, J. P. Home, G. Imreh, M. J. McDonnell, D. N. Stacey, D. J. Szwer, and A. M. Steane. A long-lived memory qubit on a low-decoherence quantum bus. *arXiv*, 0710.4421v1 [quant-ph], 2007. [cited at p. 90, 94]
- [LLC⁺09] D. R. Leibbrandt, J. Labaziewicz, R. J. Clark, I. L. Chuang, R. Epstein, C. Ospelkaus, J. Wesenberg, J. Bollinger, D. Leibfried, D. J. Wineland, D. Stick, J. Sterk, C. Monroe, C. S. Pai, Y. Low, R. Frahm, and R. E. Slusher. Demonstration of a scalable, multiplexed ion trap for quantum information processing. *Quantum Information and Computation*, 9(11):901, 2009. [cited at p. 40, 47, 73]

- [LOJ⁺05] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-lived qubit memory using atomic ions. *Physical Review Letters*, 95:060502, 2005. [cited at p. 4, 89, 90, 98]
- [LRH⁺04] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J. P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-selective photo-ionization for calcium ion trapping. *Physical Review A*, 69:012711, 2004. [cited at p. 14, 36, 80, 99, 177]
- [LW00] J.-M. Lee and K. Watkins. Laser removal of oxides and particles from copper surfaces for microelectronic fabrication. *Optics Express*, 7(2):68, 2000. [cited at p. 77]
- [MADA99] P. Meja, M. Autric, P. Delaporte, and P. Alloncle. Dry laser cleaning of anodized aluminium. *Applied Physics A*, 69:S343, 1999. [cited at p. 77, 78]
- [MBG⁺87] A. Munch, M. Berkler, Ch. Gerz, D. Wilsdorf, and G. Werth. Precise ground-state hyperfine splitting in ^{173}Yb II. *Physical Review A*, 35(10):4147, 1987. [cited at p. 90]
- [McC10] A. N. McCaughan. *High-Q Superconducting Coplanar Waveguide Resonators for Integration into Molecule Ion Traps*. PhD thesis, Massachusetts Institute of Technology, 2010. [cited at p. 174]
- [MCG⁺07] J. Majer, J. M. Chow, J. M. Gambetta, Jens Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Coupling superconducting qubits via a cavity bus. *Nature*, 449:443, 2007. [cited at p. 122]
- [MHS⁺11] D. L. Moehring, C. Highstrete, D. Stick, K. M. Fortier, R. Haltli, C. Tigges, and M. G. Blain. Design, fabrication, and experimental demonstration of junction surface ion traps. *New Journal of Physics*, 13:075018, 2011. [cited at p. 8, 61, 173]
- [MKV⁺09] T. Monz, K. Kim, A. Villar, P. Schindler, M. Chwalla, M. Riebe, C. F. Roos, H. Häffner, W. Hänsel, M. Hennrich, and R. Blatt. Realization of universal ion-trap quantum computation with decoherence-free qubits. *Physical Review Letters*, 103:200503, 2009. [cited at p. 4]
- [MMK⁺95] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Physical Review Letters*, 75:4714, 1995. [cited at p. 2, 90]
- [MMO⁺07] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe. Entanglement of single-atom quantum bits at a distance. *Nature*, 449:68, 2007. [cited at p. 6]
- [MRS⁺94] E. Matthias, M. Reichling, J. Siegel, O. W. Käding, S. Petzoldt, H. Skurk, P. Bizenberger, and E. Neske. The influence of thermal diffusion on laser ablation of metal films. *Applied Physics A*, 58:129, 1994. [cited at p. 79]
- [MS99] K. Mølmer and A. Sørensen. Multiparticle entanglement of hot trapped ions. *Physical Review Letters*, 82(9):1835, 1999. [cited at p. 110]

- [MSB⁺11] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt. 14-qubit entanglement: Creation and coherence. *Physical Review Letters*, 106:130506, 2011. [cited at p. 5]
- [MSW⁺08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, 100:200502, 2008. [cited at p. 3, 47, 56, 98]
- [MW01] F. Mintert and C. Wunderlich. Ion-trap quantum logic using long-wavelength radiation. *Physical Review Letters*, 87:257904, 2001. [cited at p. 114]
- [MWRM96] K. Mann, B. Wolff-Rottke, and F. Müller. Cleaning of optical surfaces by excimer laser radiation. *Applied Surface Science*, 96-98:463, 1996. [cited at p. 77, 78]
- [Mye07] A. H. Myerson. First year report. Technical report, University of Oxford, 2007. [cited at p. 25, 26, 29, 30]
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. [cited at p. 1, 105]
- [Obe99] H. Oberst. Resonance fluorescence of single barium ions. Master's thesis, Universität Innsbruck, 1999. [cited at p. 51]
- [OIB⁺07] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Physical Review A*, 75:042329, 2007. [cited at p. 71, 112, 178]
- [OLA⁺08] C. Ospelkaus, C. E. Langer, J. M. Amini, K. R. Brown, D. Leibfried, and D. J. Wineland. Trapped-ion quantum logic gates based on oscillating magnetic fields. *Physical Review Letters*, 101:090502, 2008. [cited at p. 89, 108, 116, 119, 178]
- [Ols72] L. Olschewski. Determination of the nuclear magnetic moments on free ⁴³Ca-, ⁸⁷Sr-, ¹³⁵Ba-, ¹³⁷Ba-, ¹⁷¹Yb- and ¹⁷³Yb-atoms by means of optical pumping. *Zeitschrift für Physik*, 249(3):205, 1972. [cited at p. 91]
- [OS11] A. Ozakin and F. Shaikh. Stability analysis of surface ion traps. *arXiv*, 1109.2160v1 [quant-ph], 2011. [cited at p. 41]
- [OWC⁺11] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476:181, 2011. [cited at p. 5, 89, 90, 95, 98, 117, 119, 124, 172]
- [OYM⁺07] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe. Manipulation and detection of a trapped Yb⁺ hyperfine qubit. *Physical Review A*, 76:052314, 2007. [cited at p. 90]
- [Phi07] C. Phipps. *Laser Ablation and its Applications*, volume 129 of *Springer Series in Optical Sciences*. Springer Berlin, 2007. [cited at p. 77]
- [Poz05] D. M. Pozar. *Microwave Engineering*. Wiley, 3rd edition, 2005. [cited at p. 115, 123, 150]
- [PRM96] S. Petzoldt, J. Reif, and E. Matthias. Laser plasma threshold of metals. *Applied Surface Science*, 96-98:199, 1996. [cited at p. 78]

- [RCK⁺06] C. F. Roos, M. Chwalla, K. Kim, M. Riebe, and R. Blatt. ‘designer atoms’ for quantum metrology. *Nature*, 443:316, 2006. [cited at p. 3]
- [RG02] E. Rubiola and V. Giordano. Advanced interferometric phase and amplitude noise measurements. *Review of Scientific Instruments*, 73(6):2445, 2002. [cited at p. 130]
- [RHR⁺04] M. Riebe, H. Häffner, C. F. Roos, W. Hänsel, J. Benhelm, G. P. T. Lancaster, T. W. Körber, C. Becher, F. Schmidt-Kaler, D. F. V. James, and R. Blatt. Deterministic quantum teleportation with atoms. *Nature*, 429:734, 2004. [cited at p. 6]
- [RLB⁺06] R. Reichle, D. Leibfried, R. B. Blakestad, J. Britton, J. D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Transport dynamics of single ions in segmented microstructured Paul trap arrays. *Fortschritte der Physik*, 54(8-10):666, 2006. [cited at p. 42]
- [RMG77] M. Rageh, D. V. Morgan, and A. E. Guile. Charge storage and the effect of ultraviolet radiation on aluminium oxide films. *Journal of Physics D: Applied Physics*, 10:2269, 1977. [cited at p. 69]
- [Rob10] J. Robin. Laser intensity stabilisation by a PID controller on a FPGA. Master’s thesis, University of Oxford, 2010. [cited at p. 28]
- [Roo08] C. F. Roos. Ion trap quantum gates with amplitude-modulated laser beams. *New Journal of Physics*, 10:013002, 2008. [cited at p. 111]
- [SBM⁺11] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt. Experimental repetitive quantum error correction. *Science*, 332:1059, 2011. [cited at p. 6]
- [Sch04] D. Schrader. *A neutral atom quantum register*. PhD thesis, University of Bonn, 2004. [cited at p. 104]
- [Sch10] R. Schmied. Electrostatics of gapped and finite surface electrodes. *New Journal of Physics*, 12:023038, 2010. [cited at p. 117]
- [SCR⁺06] S. Seidelin, J. Chiaverini, R. Reichle, J. J. Bollinger, D. Leibfried, J. Britton, J. H. Wesenberg, R. B. Blakestad, R. J. Epstein, D. B. Hume, W. M. Itano, J. D. Jost, Langer C., R. Ozeri, N. Shiga, and Wineland D. J. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Physical Review Letters*, 96:253003, 2006. [cited at p. 39, 40, 58]
- [SDM⁺11] S. Selvarajan, N. Daniilidis, S. Möller, R. Clark, F. Ziesel, K. Singer, F. Schmidt-Kaler, and H. Häffner. Electric field compensation and sensing with a single ion in a planar trap. *arXiv*, 1106.2730v1 [physics.atom-ph], 2011. [cited at p. 56]
- [SFH⁺10] D. Stick, K. M. Fortier, R. Haltli, C. Highstrete, D. L. Moehring, C. Tigges, and M. G. Blain. Demonstration of a microfabricated surface electrode ion trap. *arXiv*, 1008.0990v2 [physics.ins-det], 2010. [cited at p. 61, 62, 77, 78]
- [SGA⁺04] S. Stahl, F. Galve, J. Alonso, S. Djekic, W. Quint, T. Valenzuela, J. Verdú, M. Vogel, and G. Werth. A planar Penning trap. *The European Physical Journal D*, 32(1):139, 2004. [cited at p. 39]

- [SH92] M. A. Schilbach and A. V. Hamza. Sapphire (1120) surface: Structure and laser-induced desorption of aluminum. *Physical Review B*, 45(11):6197, 1992. [cited at p. 78]
- [SKDB10] G. Shu, N. Kurz, M. R. Dietrich, and B. B. Blinov. Efficient fluorescence collection from trapped ions with an integrated spherical mirror. *Physical Review A*, 81:042321, 2010. [cited at p. 29, 148]
- [SKK⁺00] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe. Experimental entanglement of four particles. *Nature*, 404:256, 2000. [cited at p. 107]
- [SLA⁺08] D. N. Stacey, D. M. Lucas, D. T. C. Allcock, D. J. Szwer, and S. C. Webster. Optical Bloch equations with multiply connected states. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 41:085502, 2008. [cited at p. 99]
- [SM99] A. Sørensen and K. Mølmer. Quantum computation with ions in thermal motion. *Physical Review Letters*, 8(9):1971, 1999. [cited at p. 110]
- [SM00] A. Sorensen and K. Molmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A*, 62:022311, 2000. [cited at p. 110]
- [SNJ⁺11] E. W. Streed, B. G. Norton, A. Jechow, T. J. Weinhold, and D. Kielpinski. Imaging of trapped ions with a microfabricated optic for quantum information processing. *Physical Review Letters*, 106:010502, 2011. [cited at p. 76]
- [SNRWS11] A. Safavi-Naini, P. Rabl, P. Weck, and H. R. Sadeghpour. A microscopic model of electronic field noise heating in ion traps. *arXiv*, 1106.1949v1 [quant-ph], 2011. [cited at p. 76, 84]
- [SSHM06] J. D. Sterk, D. Stick, W. K. Hensinger, and C. Monroe. Universal vacuum system for chip ion traps. In *Trapped Ion Quantum Computing Workshop (February 21-24, 2006, Boulder, CO)*, 2006. [cited at p. 40]
- [SSWH11] J. D. Siverns, L. R. Simkins, S. Weidt, and W. K. Hensinger. On the application of radio frequency voltages to ion traps via helical resonators. *arXiv*, 1106.5013v2 [physics.atom-ph], 2011. [cited at p. 63]
- [Sta09] D. N. Stacey. Magnetic dipole transitions. Technical report, University of Oxford, 2009. [cited at p. 92]
- [Ste07] A. M. Steane. How to build a 300 bit, 1 Giga-operation quantum computer. *Quantum Information and Computation*, 7(3):171, 2007. [cited at p. 7]
- [Sti07] D. Stick. *Fabrication and Characterization of Semiconductor Ion Traps for Quantum Information Processing*. PhD thesis, University of Michigan, 2007. [cited at p. 33]
- [Szw09] D. J. Szwer. *High fidelity readout and protection of a ⁴³Ca⁺ trapped ion qubit*. PhD thesis, University of Oxford, 2009. [cited at p. 4, 18, 106]
- [SZW11] N. Solmeyer, K. Zhu, and D. S. Weiss. Note: Mounting ultra-high vacuum windows with low stress-induced birefringence. *Review of Scientific Instruments*, 82:066105, 2011. [cited at p. 150]

- [TBJ⁺11] N. Timony, I. Baumgart, M. Johannng, A. F. Varon, M. B. Plenio, A. Retzker, and Ch. Wunderlich. Quantum gates and memory using microwave dressed states. *Nature*, 476:185, 2011. [cited at p. 114]
- [TH00] D. P. Taylor and H. Helvajian. Pulsed UV laser induced desorption of ions from aluminum. *Surface Science*, 451(1-3):68, 2000. [cited at p. 78]
- [TKK⁺00] Q. A. Turchette, D. Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A*, 61:063418, 2000. [cited at p. 71, 75]
- [TPR⁺03] G. Tommaseo, T. Pfeil, G. Revalde, G. Werth, P. Indelicato, and J. P. Desclaux. The g_J -factor in the ground state of Ca^+ . *The European Physical Journal D*, 25:113, 2003. [cited at p. 91]
- [Tre08] P. Treutlein. *Coherent manipulation of ultracold atoms on atom chips*. PhD thesis, Ludwig-Maximilians-Universität München, 2008. [cited at p. 116, 122]
- [UBV⁺10] H. Uys, M. J. Biercuk, A. P. VanDevender, C. Ospelkaus, D. Meiser, R. Ozeri, and J. J. Bollinger. Decoherence due to elastic Rayleigh scattering. *Physical Review Letters*, 105:200401, 2010. [cited at p. 112]
- [WD94] D. P. Woodruff and T. A. Delchar. *Modern Techniques of Surface Science*. Cambridge University Press, 2nd edition, 1994. [cited at p. 76]
- [WEL⁺07] J. H. Wesenberg, R. J. Epstein, D. Leibfried, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Fluorescence during Doppler cooling of a single trapped atom. *Physical Review A*, 76:053416, 2007. [cited at p. 58, 70]
- [Wes08] J. H. Wesenberg. Electrostatics of surface-electrode ion traps. *Physical Review A*, 78:063410, 2008. [cited at p. 40, 42]
- [WGL⁺10] S. X. Wang, Y. Ge, J. Labaziewicz, E. Dauler, K. Berggren, and I. L. Chuang. Superconducting microfabricated ion traps. *Applied Physics Letters*, 97(24):244102, 2010. [cited at p. 76]
- [WL11] D. J. Wineland and D. Leibfried. Quantum information processing and metrology with trapped ions. *Laser Physics Letters*, 8(3):175, 2011. [cited at p. 3]
- [WLG⁺09] S. X. Wang, J. Labaziewicz, Y. Ge, R. Shewmon, and I. L. Chuang. Individual addressing of ions using magnetic field gradients in a surface-electrode ion trap. *Applied Physics Letters*, 94:094103, 2009. [cited at p. 5]
- [WLL⁺11] S. X. Wang, G. H. Low, N. S. Lachenmyer, P. F. Herskind, and I. L. Chuang. Laser-induced charging of microfabricated ion traps. *arXiv*, 1108.0092v1 [cond-mat.mtrl-sci], 2011. [cited at p. 71, 98]
- [WMI⁺98a] D. J. Wineland, C. Monroe, W. M. Itano, B. E. King, D. Leibfried, D. M. Meekhof, C. Myatt, and C. Wood. Experimental primer on the trapped ion quantum computer. *Fortschritte der Physik*, 46(4-5):363, 1998. [cited at p. 47, 48, 49, 107]

BIBLIOGRAPHY

- [WMI⁺98b] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259, 1998. [cited at p. 63, 66, 75]
- [ZSM05] L. Zhu, H. Shi, and W. Menzel. Coupling behaviors of quarter-wavelength impedance transformers for wideband CPW bandpass filters. *IEEE Microwave and Wireless Components Letters*, 15(1):13, 2005. [cited at p. 123]

Appendices

A

Vacuum preparation

Days like these will be the death of me.

Butcher Boy

A.1 Vacuum compatible materials

The following table contains a list of materials present in any vacuum systems we've used that have reached pressures of $< 10^{-11}$ Torr after the preparation procedure outlined below and a vacuum bake at 190°C. For each material we give a rough estimate of the maximum surface area we've used.

Table A.1: Vacuum compatible materials

Material	Approx. quantity
Stainless steel (316LN and 304)	2500 cm ²
Silver-plated stainless steel	25 cm ²
Polyether ether ketone (PEEK)	100 cm ²
Aluminium (machined finish)	100 cm ²
Vespel SP-3	50 cm ²
Gold-plated brass	20 cm ²
Kapton-wrapped silver-plated copper wire	400 cm
Gold-plated beryllium copper	10 cm ²
Epo-Tek H20E conductive epoxy	1 cm ²
Epo-Tek 353ND epoxy	1 cm ²

A.1.1 Preparation

All parts apart from the epoxy are cleaned first by submersion in Decon 90 detergent. They are then rinsed in de-ionised (DI) water, submersed in acetone and then given a final rinse in DI water. After each stage, any residual liquid is blown off with clean compressed gas rather than being allowed to dry. Unless the parts are fragile (e.g. viewports) an ultrasonic bath is used to agitate the part during each cleaning stage. Parts which are fabricated out of stainless steel only are given an air bake for a few hours at $> 350^{\circ}\text{C}$ to promote oxide growth on the surface in order to reduce hydrogen outgassing.

Mixing the 353ND epoxy incorporates air bubbles into it which could cause problems in UHV work. These are removed before use by pouring the epoxy into a wide dish and heating to 40°C .

B

$^{43}\text{Ca}^+$ $S_{1/2}$ states and M1 matrix elements at intermediate field

B.1 $S_{1/2}$ hyperfine states

The projection of the angular momentum onto the external field M remains a good quantum number at all fields and $J = 1/2$ so any state can be a superposition of at most two basis states:

$$\begin{aligned} |n(M_n)\rangle &= \alpha_n |JI, M_J = -\frac{1}{2}, M_I = M_n + \frac{1}{2}\rangle \\ &\quad + \beta_n |JI, M_J = \frac{1}{2}, M_I = M_n - \frac{1}{2}\rangle \\ &\equiv \alpha_n |-\frac{1}{2}, M_n + \frac{1}{2}\rangle + \beta_n |\frac{1}{2}, M_n - \frac{1}{2}\rangle. \end{aligned} \tag{B.1}$$

Table B.1: The basis state coefficients for each M state at 23 G.

M	Basis state				Basis state			
	M_J	M_I	α_m	α_n	M_J	M_I	β_m	β_n
-4	-1/2	-7/2	1	0				
-3	-1/2	-5/2	0.9377	0.3475	1/2	-7/2	0.3475	-0.9377
-2	-1/2	-3/2	0.8703	0.4926	1/2	-5/2	0.4926	-0.8703
-1	-1/2	-1/2	0.7964	0.6047	1/2	-3/2	0.6047	-0.7964
0	-1/2	1/2	0.7141	0.7000	1/2	-1/2	0.7000	-0.7141
1	-1/2	3/2	0.6200	0.7846	1/2	1/2	0.7846	-0.6200
2	-1/2	5/2	0.5075	0.8616	1/2	3/2	0.8616	-0.5075
3	-1/2	7/2	0.3598	0.9330	1/2	5/2	0.9330	-0.3598
4					1/2	7/2	1	0

Table B.2: The basis state coefficients for each M state at 146.094 G.

M	Basis state				Basis state			
	M_J	M_I	α_m	α_n	M_J	M_I	β_m	β_n
-4	-1/2	-7/2	1	0				
-3	-1/2	-5/2	0.9483	0.3175	1/2	-7/2	0.3175	-0.9483
-2	-1/2	-3/2	0.8906	0.4548	1/2	-5/2	0.4548	-0.8906
-1	-1/2	-1/2	0.8255	0.5645	1/2	-3/2	0.5645	-0.8255
0	-1/2	1/2	0.7503	0.6611	1/2	-1/2	0.6611	-0.7503
1	-1/2	3/2	0.6610	0.7504	1/2	1/2	0.7504	-0.6610
2	-1/2	5/2	0.5497	0.8354	1/2	3/2	0.8354	-0.5497
3	-1/2	7/2	0.3964	0.9181	1/2	5/2	0.9181	-0.3964
4					1/2	7/2	1	0

Table B.3: The basis state coefficients for each M state at 287.783 G.

M	Basis state				Basis state			
	M_J	M_I	α_m	α_n	M_J	M_I	β_m	β_n
-4	-1/2	-7/2	1	0				
-3	-1/2	-5/2	0.9576	0.2880	1/2	-7/2	0.2880	-0.9576
-2	-1/2	-3/2	0.9096	0.4155	1/2	-5/2	0.4155	-0.9096
-1	-1/2	-1/2	0.8541	0.5202	1/2	-3/2	0.5202	-0.8541
0	-1/2	1/2	0.7882	0.6154	1/2	-1/2	0.6154	-0.7882
1	-1/2	3/2	0.7071	0.7071	1/2	1/2	0.7071	-0.7071
2	-1/2	5/2	0.6011	0.7992	1/2	3/2	0.7992	-0.6011
3	-1/2	7/2	0.4455	0.8953	1/2	5/2	0.8953	-0.4455
4					1/2	7/2	1	0

B.2 M1 transition frequencies and matrix elements

Table B.4: Magnetic dipole transitions at a field of 23 G. The qubit transition is in bold.

	M_m	M_n	Frequency (GHz)	df/dB_0 (MHz/G)	Matrix element ($\mu_B B_{\mu w}$)
a	-4	-3	3.282146	2.464	1.3276
b	-3	-3	3.274222	2.126	0.6524
c	-3	-2	3.266273	1.784	1.1554
d	-2	-3	3.266260	1.784	0.2423
e	-2	-2	3.258311	1.442	0.8583
f	-2	-1	3.250322	1.098	0.9813
g	-1	-2	3.250309	1.097	0.4217
h	-1	-1	3.242321	0.7526	0.9643
i	-1	0	3.234293	0.4046	0.8052
j	0	-1	3.234280	0.4040	0.5993
k	0	0	3.226252077	0.05598 (1.216 kHz/G²)	1.0009
l	0	+1	3.218184	-0.2956	0.6269
m	+1	0	3.218171	-0.2961	0.7776
n	+1	+1	3.210103	-0.6477	0.9740
o	+1	+2	3.201994	-1.003	0.4455
p	+2	+1	3.201981	-1.003	0.9571
q	+2	+2	3.193872	-1.358	0.8756
r	+2	+3	3.185721	-1.717	0.2586
s	+3	+2	3.185708	-1.718	1.1382
t	+3	+3	3.177558	-2.077	0.6722
u	+4	+3	3.169352	-2.440	1.3210

Table B.5: Magnetic dipole transitions at a field of 146.094 G. The qubit transition is in bold.

	M_m	M_n	Frequency (GHz)	df/dB_0 (MHz/G)	Matrix element ($\mu_B B_{\mu\nu}$)
a	-4	-3	3.589033	2.520	1.3425
b	-3	-3	3.543001	2.237	0.6029
c	-3	-2	3.495826	1.940	1.1957
d	-2	-3	3.495742	1.940	0.2044
e	-2	-2	3.448567	1.643	0.8109
f	-2	-1	3.400062	1.330	1.0408
g	-1	-2	3.399978	1.330	0.3634
h	-1	-1	3.351474	1.016	0.9329
i	-1	0	3.301520	0.6849	0.8769
j	0	-1	3.301436	0.6844	0.5283
k	0	0	3.251482	0.3528	0.9932
l	0	+1	3.199941078	0 (1.208 kHz/G²)	0.7022
m	+1	0	3.199857	-0.00057	0.7023
n	+1	+1	3.148316	-0.3534	0.9931
o	+1	+2	3.095027	-0.7307	0.5145
p	+2	+1	3.094943	-0.7313	0.8875
q	+2	+2	3.041654	-1.109	0.9194
r	+2	+3	2.986425	-1.515	0.3085
s	+3	+2	2.986341	-1.515	1.0858
t	+3	+3	2.931112	-1.921	0.7287
u	+4	+3	2.873631	-2.362	1.2998

Table B.6: Magnetic dipole transitions at a field of 287.783 G. The qubit transition is in bold.

	M_m	M_n	Frequency (GHz)	df/dB_0 (MHz/G)	Matrix element ($\mu_B B_{\mu\nu}$)
a	-4	-3	3.949772	2.570	1.3558
b	-3	-3	3.867368	2.337	0.5523
c	-3	-2	3.781466	2.086	1.2332
d	-2	-3	3.781301	2.085	0.1695
e	-2	-2	3.695399	1.834	0.7568
f	-2	-1	3.605294	1.560	1.0998
g	-1	-2	3.605129	1.560	0.3060
h	-1	-1	3.515025	1.286	0.8895
i	-1	0	3.420035	0.9829	0.9531
j	0	-1	3.419871	0.9824	0.4532
k	0	0	3.324881	0.6796	0.9712
l	0	+1	3.224114	0.3401	0.7891
m	+1	0	3.223949	0.3395	0.6161
n	+1	+1	3.123181794	0 (1.257 kHz/G²)	1.0011
o	+1	+2	3.015435	-0.3883	0.6018
p	+2	+1	3.015270	-0.3889	0.8001
q	+2	+2	2.907524	-0.7772	0.9619
r	+2	+3	2.791111	-1.233	0.3792
s	+3	+2	2.790946	-1.234	1.0130
t	+3	+3	2.674533	-1.690	0.7986
u	+4	+3	2.546787	-2.246	1.2675

List of abbreviations

Symbol	Description
ADC	analogue-to-digital converter
AM	amplitude modulation
AOM	acousto-optic modulator
AR	anti-reflection
ASE	amplified spontaneous emission
AWG	arbitrary waveform generator
BNC	bayonet Neill-Concelman
BSE	back-scattered electron
CCD	charge coupled device
CF	ConFlat
CMOS	complementary metal-oxide-semiconductor
CPGA	ceramic pin grid array
CPW	coplanar waveguide
CVD	chemical vapour deposition
CW	continuous wave
D3	the 3rd David
DAC	digital-to-analogue convertor
DDS	direct digital synthesis
DI	de-ionised
EMCCD	electron multiplying charge-coupled device
EOM	electro-optic modulator
ESL	equivalent series inductance
FC-APC	fixed connection - angled physical contact
FPGA	field-programmable gate array
FSR	free spectral range
GHZ	Greenberger-Horne-Zeilinger
GPIB	general purpose interface bus
IPA	isopropyl alcohol
IL	insertion loss
ITO	indium tin oxide
LCU	laser control unit
LO	local oscillator
MEMS	micro-electro-mechanical system
MWO	Microwave Office
NA	numerical aperture
Nd:YAG	neodymium-doped yttrium aluminium garnet
NF	noise figure
NIM	nuclear instrumentation module
NIST	National Institute of Standards and Technology

LIST OF ABBREVIATIONS

Symbol	Description
NPBS	non-polarising beamsplitter
NPL	UK National Physical Laboratory
PCB	printed circuit board
PBS	polarising beamsplitter
PCI	peripheral component interconnect
PDH	Pound Drever Hall
PEEK	polyetheretherketone
PI	photoionisation
PID	proportional-integral-derivative
PMT	photomultiplier tube
PTFE	polytetrafluorethylene
QIP	quantum information processing
QND	quantum non-demolition
SMA	subminiature version A
SRF	self-resonant frequency
TAC	time-to-amplitude converter
TEOS	tetraethyl orthosilicate
TTL	transistor-transistor logic
UHV	ultra-high vacuum
VNA	vector network analyser
VSWR	voltage standing wave ratio